




FORTINET®



FortiOS™ Handbook

VERSION 6.0.2

**FORTIOS
VERSION
6.0**

FORTINET DOCUMENT LIBRARY

<https://docs.fortinet.com>

FORTINET VIDEO GUIDE

<https://video.fortinet.com>

FORTINET KNOWLEDGE BASE

<http://kb.fortinet.com>

FORTINET BLOG

<https://blog.fortinet.com>

CUSTOMER SERVICE & SUPPORT

<https://support.fortinet.com>

<http://cookbook.fortinet.com/how-to-work-with-fortinet-support/>

FORTINET COOKBOOK

<http://cookbook.fortinet.com>

FORTINET TRAINING AND CERTIFICATION PROGRAM

<https://www.fortinet.com/support-and-training/training.html>

NSE INSTITUTE

<https://training.fortinet.com/>

FORTIGUARD CENTER

<https://fortiguard.com>

FORTICAST

<http://forticast.fortinet.com>

END USER LICENSE AGREEMENT

<https://www.fortinet.com/doc/legal/EULA.pdf>



July 26, 2018

FortiOS™ Handbook

01-602-481047-20180726

TABLE OF CONTENTS

Chapter 1 - What's New for FortiOS 6.0	48
Fortinet Security Fabric	48
Security Fabric automation	48
Security rating	48
Solution and service integration	48
Multi-cloud support (Security Fabric connectors)	50
Manageability	51
Asset tagging	51
FortiSwitch network assisted device detection and destination name resolution	51
Global security profiles	51
Networking	52
SD-WAN improvements	52
Cloud-assisted one-click VPN	53
IPv6 enhancements	53
NAT enhancements	53
EMAC-VLAN support	53
Security	54
FortiGuard virus outbreak prevention	54
FortiGuard content disarm and reconstruction	54
Application groups for NGFW policies	54
Application control rule sequencing	54
Threat Feeds (external dynamic block lists)	55
FortiAP-S bridge mode security profiles	55
Chapter 2 - Getting Started	56
Differences between models	56
What's new in FortiOS 6.0	57
Installation	58
Quick installation using DHCP	58
Installing a FortiGate in NAT/route mode	59
Using a virtual wire pair	61
Using the GUI	63
Connecting to the GUI using a web browser	63
Menus	64
Dashboard	66

Feature Visibility.....	69
Tables.....	70
Text strings.....	72
Using the CLI.....	74
Connecting to the CLI.....	74
Command syntax.....	79
Sub-commands.....	84
Permissions.....	88
Tips.....	89
FortiExplorer for iOS.....	97
Getting started with FortiExplorer.....	97
LED specifications.....	102
Sample FortiGate faceplates.....	102
LED status codes.....	103
About alarm levels.....	104
LED status codes for ports.....	104
Inspection mode.....	106
Basic administration.....	110
Registration.....	110
System Settings.....	110
Passwords.....	114
Firmware.....	116
Configuration backups.....	124
FortiGuard.....	130
FortiCloud.....	138
Troubleshooting your FortiGate installation.....	140
Resources.....	143
Best Practices.....	143
The Fortinet Cookbook.....	143
The Fortinet Video Library.....	143
The FortiOS Handbook.....	143
Fortinet Support.....	143
Chapter 3 - Authentication.....	144
What's new in FortiOS 6.0.1.....	145
What's new in FortiOS 6.0.....	145
Introduction to authentication.....	146
What is authentication?.....	146
Methods of authentication.....	146
Types of authentication.....	148
Single sign-on authentication for users.....	151
User's view of authentication.....	152
FortiGate administrator's view of authentication.....	153

General authentication settings.....	154
Authentication servers.....	155
FortiAuthenticator servers.....	155
RADIUS servers.....	155
LDAP servers.....	163
TACACS+ servers.....	172
POP3 servers.....	174
SSO servers.....	174
RSA ACE (SecurID) servers.....	176
Users and user groups.....	181
Users.....	181
User groups.....	196
Managing guest access.....	202
Configuring guest user access.....	202
Guest access in a retail environment.....	204
Configuring authenticated access.....	208
Authentication timeout.....	208
Password policy.....	209
Authentication protocols.....	211
Authentication in captive portals.....	212
Authentication in security policies.....	212
Authentication replacement messages.....	213
Kerberos authentication for explicit web and transparent web proxy users.....	218
VPN authentication.....	228
Captive portals.....	233
Introduction to captive portals.....	233
Configuring a captive portal.....	233
Customizing captive portal pages.....	235
Certificate-based authentication.....	241
What is a security certificate?.....	241
Certificates overview.....	242
Managing X.509 certificates.....	245
Troubleshooting certificates.....	250
Configuring certificate-based authentication.....	255
Support for per-VDOM certificates.....	257
Example — Generate a CSR on the FortiGate unit.....	258
Example — Generate and Import CA certificate with private key pair on OpenSSL.....	259
Example — Generate an SSL certificate in OpenSSL.....	260
Single sign-on using a FortiAuthenticator unit.....	263
User's view of FortiAuthenticator SSO authentication.....	263
Administrator's view of FortiAuthenticator SSO authentication.....	263
Configuring the FortiAuthenticator unit.....	264

Configuring the FortiGate unit	265
Configuring the FortiClient SSO Mobility Agent	266
Viewing SSO authentication events on the FortiGate unit	266
Single sign-on to Windows AD	268
Introduction to SSO with Windows AD	268
Configuring SSO to Windows AD	268
FortiOS FSSO log messages	273
Testing FSSO	274
Troubleshooting FSSO	274
Agent-based FSSO	276
Introduction to agent-based FSSO	276
FSSO NTLM authentication support	282
Agent installation	285
Configuring the FSSO collector agent for Windows AD	290
Configuring the FSSO TS agent for Citrix	302
Configuring FSSO with Novell networks	303
Configuring FSSO advanced settings	305
Configuring FSSO on FortiGate units	310
FortiOS FSSO log messages	315
Testing FSSO	317
Troubleshooting FSSO	318
SSO using RADIUS accounting records	321
User's view of RADIUS SSO authentication	321
Configuration overview	321
Configuring the RADIUS server	321
Creating the FortiGate RADIUS SSO agent	322
Defining local user groups for RADIUS SSO	324
Creating security policies	325
Example - webfiltering for student and teacher accounts	326
Monitoring authenticated users	329
Monitoring firewall users	329
Monitoring SSL VPN users	329
Monitoring IPsec VPN users	330
Monitoring users quarantine	330
Examples and troubleshooting	332
Firewall authentication example	332
LDAP dial-in using member-attribute example	338
RADIUS SSO example	340
Troubleshooting	348
Chapter 4 - Best Practices	349
Introduction	349
General considerations	349

Customer service and technical support	349
Fortinet Knowledge Base	349
System and performance	350
Performance	350
Shutting down	350
Migration	351
Information gathering	351
Object and policy migration	351
Testing and validation	351
Going live and obtaining feedback	352
Adding new services	352
Environmental specifications	353
Grounding	353
Rack mounting	353
Firmware	355
Firmware change management	355
Performing a firmware upgrade	358
Performing a firmware downgrade	359
Performing a configuration backup	359
Security Profiles (AV, Web Filtering etc.)	362
Firewall	362
Security	362
Authentication	363
Antivirus	363
Antispam	363
Intrusion Prevention System (IPS)	363
Email filter	364
URL filtering	364
Web filtering	365
Patch management	365
Policy configuration	365
Networking	367
Routing configuration	367
Advanced routing	367
Network Address Translation (NAT)	368
Transparent Mode	368
Using virtual IPs (VIPs)	368
FGCP high availability	369
Heartbeat interfaces	369
Interface monitoring (port monitoring)	370
WAN Optimization	371
Virtual Domains (VDOMs)	373

Per-VDOM resource settings	373
Virtual domains in NAT/Route mode	373
Virtual clustering	373
Explicit proxy	374
Wireless	375
Encryption and authentication	375
Geographic location	375
Network planning	375
Lowering the power level to reduce RF interference	376
Wireless client load balancing	376
Local bridging	376
Advertising SSIDs	377
Using static IPs in a CAPWAP configuration	377
Logging and reporting	378
Log management	378
System memory and hard disks	379
Chapter 5 - FortiOS Carrier	380
What's new in FortiOS Carrier 6.0	380
Overview of FortiOS Carrier features	382
Overview	382
MMS	382
GTP	382
MMS Concepts	382
MMS background	382
How FortiOS Carrier processes MMS messages	386
MMS protection profiles	396
Bypassing MMS protection profile filtering based on carrier endpoints	397
Applying MMS protection profiles to MMS traffic	397
MMS Configuration	397
MMS profiles	397
MMS Content Checksum	412
Notification List	413
Message Flood	415
Duplicate Message	417
Carrier Endpoint Filter Lists	418
Message flood protection	421
Duplicate message protection	432
Employing MMS Security features	439
GTP basic concepts	453
PDP Context	454
GPRS security	455
Parts of a GTPv1 network	456

Radio access	457
Transport	457
Billing and records	460
GPRS network common interfaces	461
GTP Configuration	462
Introduction to GTP	462
GTP Profile	464
Configuring GTP on FortiOS Carrier	484
GTP message type filtering	493
GTP identity filtering	500
SCTP Concepts	506
SCTP Firewall	508
Troubleshooting	510
FortiOS Carrier diagnose commands	510
Applying IPS signatures to IP packets within GTP-U tunnels	511
GTP packets are not moving along your network	512
Chapter 6 - Firewall	517
Fundamentals	517
FortiGate firewall components	517
Firewall optimization	518
How does a FortiGate protect your network?	519
What's new for Firewall in 6.0.1	520
What's new for Firewall in 6.0	520
Firewall concepts	522
What is a firewall?	522
FortiGate modes	525
How packets are handled by FortiOS	526
Interfaces and zones	527
Access control lists	528
Firewall policies	529
Policy modes	542
Security profiles	545
NAT	559
IP pools	572
Services and TCP ports	575
VPN policies	601
DSRL	601
Interface policies	602
DoS protection	602
Local-In policies	608
Security policy 0	609
DNS traffic in NGFW policy-mode	610

Deny policies.....	610
Accept policies.....	610
Fixed port.....	610
Endpoint security.....	611
Traffic logging.....	611
IPv6.....	613
IPv6 addressing.....	613
IPv6 packet structure.....	614
Benefits of IPv6.....	614
IPv6 in FortiOS.....	614
IPv6 features.....	615
IPv6 configuration.....	637
Network defense.....	665
Monitoring.....	665
Blocking external probes.....	665
Defending against DoS attacks.....	670
Inside FortiOS: Denial of Service (DoS) protection.....	675
About DoS and DDoS attacks.....	675
FortiOS DoS and DDoS protection.....	675
FortiOS DDoS prevention.....	676
Configuration options.....	677
DoS policies.....	678
Hardware acceleration.....	678
The FortiGuard Center.....	678
Policy configuration.....	679
Viewing firewall policies.....	679
Policy names.....	681
IPv4 policy.....	682
IPv6 policy.....	687
NAT64 policy.....	689
NAT46 policy.....	691
Central SNAT.....	692
IPv4 access control list.....	695
IPv6 access control list.....	696
IPv4 DoS policy.....	697
IPv6 DoS policy.....	702
Multicast policy.....	703
SSL mirroring for policies.....	704
Object configuration.....	705
UUID support.....	705
Addresses.....	705
Virtual IPs.....	728

Configuring IP pools.....	736
Services.....	742
Firewall schedules.....	749
Secure web gateway, WAN optimization, web caching and WCCP.....	756
Before you begin.....	756
FortiGate models that support WAN optimization.....	756
Distributing WAN optimization, explicit proxy, and web caching to multiple CPU cores.....	756
Dispatching traffic to WAD worker based on source affinity.....	756
Toggling disk usage for logging or wan-opt.....	757
Example topologies relevant to WAN optimization.....	759
Basic WAN optimization topology.....	759
Out-of-path WAN optimization topology.....	760
Topology for multiple networks.....	761
WAN optimization with web caching.....	761
Explicit web proxy topologies.....	762
Explicit FTP proxy topologies.....	762
Web caching topologies.....	763
WCCP topologies.....	764
Inside FortiOS: WAN optimization.....	765
Centralize without compromising your WAN performance.....	765
FortiOS WAN optimization.....	765
Protocol optimization.....	765
Web caching.....	766
Byte caching.....	766
Server monitoring and management.....	767
SSL acceleration.....	767
VPN replacement.....	767
Road warriors and home workers.....	767
WAN optimization concepts.....	769
Client/server architecture.....	769
WAN optimization peers.....	770
Protocol optimization.....	770
Protocol optimization and MAPI.....	771
Byte caching.....	771
WAN optimization transparent mode.....	772
Operating modes and VDOMs.....	773
WAN optimization tunnels.....	773
WAN optimization and user and device identity policies, load balancing and traffic shaping.....	774
WAN optimization and HA.....	775
WAN optimization, web caching and memory usage.....	775
WAN optimization configuration.....	776

Manual (peer-to-peer) and active-passive WAN optimization.....	776
WAN optimization profiles.....	778
Monitoring WAN optimization performance.....	781
WAN optimization configuration summary.....	782
WANopt storage.....	785
WANopt cache service.....	786
Video caching.....	787
Best practices.....	789
Example basic manual (peer-to-peer) WAN optimization configuration.....	789
Example active-passive WAN optimization.....	796
Example adding secure tunneling to an active-passive WAN optimization configuration.....	803
Peers and authentication groups.....	811
Basic WAN optimization peer requirements.....	811
How FortiGate units process tunnel requests for peer authentication.....	811
Configuring peers.....	812
Configuring authentication groups.....	813
Secure tunneling.....	815
Monitoring WAN optimization peer performance.....	816
Web cache concepts.....	817
Turning on web caching for HTTP and HTTPS traffic.....	817
Turning on web caching for HTTPS traffic.....	818
Changing the ports on which to look for HTTP and HTTPS traffic to cache.....	820
Web caching and HA.....	821
Web caching and memory usage.....	821
Changing web cache settings.....	821
Web cache configuration.....	825
Forwarding URLs to forwarding servers and exempting web sites from web caching... ..	825
Monitoring web caching performance.....	827
Example web caching of HTTP and HTTPS Internet content for users on an internal network.....	827
Example reverse proxy web caching and SSL offloading for an Internet web server using a static one-to-one virtual IP.....	830
Using a FortiCache as a cache service.....	834
WCCP concepts.....	835
WCCP Cisco to FortiGate client using L2-forwarding tunneling.....	835
WCCP configuration.....	836
WCCP configuration overview.....	836
WCCP service groups, service numbers, service IDs and well known services.....	836
Example caching HTTP sessions on port 80 using WCCP.....	840
Example caching HTTP sessions on port 80 and HTTPS sessions on port 443 using WCCP.....	842
WCCP packet flow.....	845

Configuring the forward and return methods and adding authentication	845
WCCP messages	846
Troubleshooting WCCP	846
Web proxy concepts	848
Proxy policy	848
Proxy authentication	848
Proxy addresses	852
Web proxy firewall services and service groups	852
Learn client IP	853
Web proxy configuration	855
General web proxy configuration steps	855
Logging options in web proxy profiles	859
Policy matching based on referrer headers and query strings	860
Multiple web proxy PAC files in one VDOM	862
Explicit proxy concepts	864
The FortiGate explicit web proxy	864
Other explicit web proxy options	865
Proxy chaining (web proxy forwarding servers)	867
Security profiles, threat weight, device identification, and the explicit web proxy	870
Explicit web proxy sessions and user limits	871
Explicit proxy configuration	873
Configuring an external IP address for the IPv4 explicit web proxy	873
Configuring an external IP address for the IPv6 explicit web proxy	873
Restricting the IP address of the IPv4 explicit web proxy	873
Restricting the outgoing source IP address of the IPv4 explicit web proxy	873
Restricting the IP address of the explicit IPv6 web proxy	874
Restricting the outgoing source IP address of the IPv6 explicit web proxy	874
Explicit proxy firewall address types	874
Proxy auto-config (PAC) configuration	874
Unknown HTTP version	875
Authentication realm	875
Implementing botnet features	875
Adding disclaimer messages to explicit proxy policies	876
Changing HTTP headers	876
Preventing the explicit web proxy from changing source addresses	877
Kerberos and NTLM authentication	882
Kerberos authentication for explicit proxy users	882
Transparent proxy concepts	892
More about the transparent proxy	892
Transparent proxy configuration	895
FTP proxy concepts	899
The FortiGate explicit FTP proxy	899

How to use the explicit FTP proxy to connect to an FTP server.....	900
Security profiles, threat weight, device identification, and the explicit FTP proxy.....	902
Explicit FTP proxy sessions and user limits.....	902
FTP proxy configuration.....	903
General explicit FTP proxy configuration steps.....	903
Restricting the IP address of the explicit FTP proxy.....	905
Restricting the outgoing source IP address of the explicit FTP proxy.....	905
Example users on an internal network connecting to FTP servers on the Internet through the explicit FTP with RADIUS authentication and virus scanning.....	906
Diagnose commands for WAN optimization.....	911
get test {wad wccpd} <test_level>.....	911
diagnose wad.....	912
diagnose wacs.....	916
diagnose wadbd.....	917
diagnose debug application {wad wccpd} [<debug_level>].....	917
diagnose test application wad 2200.....	917
Chapter 7 - FortiView.....	920
What's new in FortiOS 6.0.....	921
What's New in FortiOS 6.0.....	921
Purpose.....	923
Overview.....	924
Enabling FortiView.....	925
FortiView feature support - platform matrix.....	925
Configuration dependencies.....	929
FortiView interface.....	931
FortiView consoles.....	935
Sources.....	936
Destinations.....	936
Applications.....	937
Cloud Applications.....	937
Web Sites.....	938
Threats.....	939
WiFi Clients.....	940
Traffic Shaping.....	940
System Events.....	940
VPN.....	941
Endpoint Vulnerability.....	942
Threat Map.....	942
Policies.....	943
Interfaces.....	944
FortiSandbox.....	944
All Sessions.....	945

Reference.....	946
Filtering options.....	947
Drill-Down Options.....	949
Columns displayed.....	950
Risk level indicators.....	956
Troubleshooting FortiView.....	958
No logging data is displayed.....	958
Logging is enabled, but data is not appearing.....	958
Chapter 8 - Fortinet Communication Ports and Protocols.....	959
What's new in FortiOS 6.0.....	959
Introduction.....	960
FortiGate open ports.....	961
FortiAnalyzer open ports.....	965
FortiAP-S open ports.....	967
FortiAuthenticator open ports.....	969
FortiClient open ports.....	972
FortiClient.....	972
FortiClient EMS.....	974
FortiClient for Chromebook.....	975
FortiClient EMS for Chromebook.....	976
FortiCloud open ports.....	977
FortiDB open ports.....	978
FortiGuard open ports.....	979
FortiMail open ports.....	982
FortiManager open ports.....	987
FortiPortal open ports.....	990
FortiSandbox open ports.....	992
Services and port numbers required for FortiSandbox.....	993
3rd-party servers open ports.....	994
Fortinet proprietary protocols.....	996
FGCP - FortiGate Clustering Protocol.....	997
FGSP - FortiGate Session Life Support Protocol.....	1004
FGFM - FortiGate to FortiManager Protocol.....	1009
SLBC - Session-aware Load Balancing Cluster.....	1012
Fortinet Security Fabric.....	1020
FortiTelemetry/On-Net/FortiClient Endpoint Compliance.....	1022
FortiGuard.....	1025
FortiLink.....	1029
FortiOS WAN optimization.....	1035
FSSO - Fortinet Single Sign-On.....	1039
OFTP - Optimized Fabric Transfer Protocol.....	1043
FortiClient EMS - Enterprise Management Server.....	1044

Chapter 9 - FortiWiFi and FortiAP Configuration Guide	1046
What's new in FortiOS 6.0.1	1046
What's new in FortiOS 6.0	1047
Introduction to wireless networking	1048
Wireless concepts	1048
Security	1049
Authentication	1051
Wireless networking equipment	1051
Automatic Radio Resource Provisioning	1052
Captive portals	1054
Introduction to captive portals	1054
Configuring a captive portal	1054
Customizing captive portal pages	1057
Configuration example - captive portal WiFi access control	1062
Configuring a WiFi LAN	1065
Overview of WiFi controller configuration	1065
Setting your geographic location	1067
Creating a FortiAP profile	1068
Defining a wireless network interface (SSID)	1070
Defining SSID groups	1079
Dynamic user VLAN assignment	1080
Configuring user authentication	1083
Configuring firewall policies for the SSID	1085
Configuring the built-in access point on a FortiWiFi unit	1087
Enforcing UTM policies on a local bridge SSID for managed smart APs	1087
Access point deployment	1089
Overview	1089
Network topology for managed APs	1089
Discovering and authorizing APs	1093
Advanced WiFi controller discovery	1099
Wireless client load balancing for high-density deployments	1101
FortiAP groups	1103
LAN port options	1103
Preventing IP fragmentation of packets in CAPWAP tunnels	1106
LED options	1107
CAPWAP bandwidth formula	1108
Wireless mesh	1110
Overview of wireless mesh	1110
Configuring a meshed WiFi network	1113
Configuring a point-to-point bridge	1116
Hotspot 2.0	1118
Combining WiFi and wired networks with a software switch	1122

Combining WiFi and wired networks with a software switch	1122
FortiAP local bridging (Private cloud-managed AP)	1124
Using bridged FortiAPs to increase scalability	1126
Using remote WLAN FortiAPs	1128
Configuring the FortiGate for remote FortiAPs	1128
Configuring the FortiAP units	1129
Preauthorizing FortiAP units	1130
Features for high-density deployments	1131
Multiple FortiAP firmware upgrades at once	1131
Power save feature	1131
Broadcast packet suppression	1132
Multicast to unicast conversion	1133
Ignore weak or distant clients	1133
Turn off the 802.11b protocol	1133
Disable low data rates	1134
Limit power	1134
Use frequency band load-balancing	1134
AP load balancing	1135
Application rate-limiting	1136
AP group management and dynamic VLAN assignment	1136
Sharing tunnel SSIDs within a single managed AP between VDOMs as a virtual AP for multi-tenancy	1136
Manual quarantine of devices on FortiAP (tunnel mode)	1136
Locate a FortiAP with LED blinking	1138
Wireless controller optimization for large deployment - AP image upgrade	1138
Control message off-loading and aeroscout enhancement	1139
Protecting the WiFi network	1140
Wireless IDS	1140
WiFi data channel encryption	1141
Protected Management Frames and Opportunistic Key Caching support	1142
Bluetooth Low Energy (BLE) Scan	1142
Preventing local bridge traffic from reaching the LAN	1143
FortiAP-S bridge mode security profiles	1144
DHCP snooping and option 82 (circuit-id) options for wireless access points	1144
Wireless network monitoring	1145
Monitoring wireless clients	1145
Monitoring rogue APs	1145
Suppressing rogue APs	1150
Monitoring wireless network health	1150
Configuring wireless network clients	1152
Windows XP client	1152
Windows 7 client	1157

Mac OS client	1161
Linux client	1163
Troubleshooting	1165
Wireless network examples	1168
Basic wireless network	1168
A more complex example	1173
Managing a FortiAP with FortiCloud	1183
FortiCloud-managed FortiAP WiFi	1183
FortiCloud-managed FortiAP WiFi without a key	1184
Using a FortiWiFi unit as a client	1186
Use of client mode	1186
Configuring client mode	1188
Support for location-based services	1190
Overview	1190
Configuring location tracking	1190
Viewing device location data on the FortiGate unit	1191
Troubleshooting	1193
FortiAP shell command through CAPWAP control tunnel	1193
Signal strength issues	1193
Throughput issues	1197
Connection issues	1198
General problems	1203
Packet sniffer	1206
Useful debugging commands	1210
Reference	1212
FortiAP web-based manager	1213
Wireless radio channels	1215
WiFi event types	1217
FortiAP CLI	1217
Chapter 10 - Hardening your FortiGate	1224
Building security into FortiOS	1225
Boot PROM and BIOS security	1225
FortiOS kernel and user processes	1225
Administration access security	1225
Network security	1227
FIPS and Common Criteria	1228
PSIRT advisories	1228
FortiOS ports and protocols	1229
FortiOS open ports	1229
Closing open ports	1229
Security best practices	1230
Install the FortiGate unit in a physically secure location	1230

Register your product with Fortinet Support	1230
Keep your FortiOS firmware up to date.....	1230
System administrator best practices.....	1230
Global commands for stronger and more secure encryption.....	1234
Disable sending malware statistics to FortiGuard.....	1235
Disable sending Security Rating statistics to FortiGuard.....	1235
Disable auto USB installation.....	1235
Set system time by synchronizing with an NTP server.....	1236
Disable the maintainer admin account.....	1236
Enable password policies.....	1236
Configure auditing and logging.....	1236
Disable unused interfaces.....	1237
Disable unused protocols on interfaces.....	1237
Use local-in policies to close open ports or restrict access.....	1238
Chapter 11 - Hardware Acceleration.....	1240
What's new in FortiOS 6.0.2.....	1240
What's new in FortiOS 6.0.....	1240
Hardware acceleration overview.....	1241
Content processors (CP4, CP5, CP6, CP8, and CP9).....	1241
Security processors (SPs).....	1244
Network processors (NP1, NP2, NP3, NP4, NP4Lite, NP6 and NP6Lite).....	1246
Enabling strict protocol header checking disables all hardware acceleration.....	1250
sFlow and NetFlow and hardware acceleration.....	1250
Checking that traffic is offloaded by NP processors.....	1250
Dedicated management CPU.....	1252
Offloading flow-based content inspection with NTurbo and IPSA.....	1253
Preventing packet ordering problems with NP4, NP6 and NP6lite FortiGates under heavy load.....	1254
NP6 and NP6lite acceleration.....	1255
NP6 session fast path requirements.....	1256
NP6Lite processors.....	1257
NP6 and NP6Lite processors and sFlow and NetFlow.....	1257
NP6 processors and traffic shaping.....	1257
NP Direct.....	1258
Viewing your FortiGate NP6 processor configuration.....	1258
Disabling NP6 and NP6lite hardware acceleration (fastpath).....	1259
Optimizing NP6 performance by distributing traffic to XAUI links.....	1259
Enabling bandwidth control between the ISF and NP6 XAUI ports.....	1261
Increasing NP6 offloading capacity using link aggregation groups (LAGs).....	1261
Configuring inter-VDOM link acceleration with NP6 processors.....	1262
Disabling offloading IPsec Diffie-Hellman key exchange.....	1265
Configuring individual NP6 processors.....	1265

Enabling per-session accounting for offloaded NP6 and NP6lite sessions	1271
Configuring NP6 session timeouts	1272
Configure the number of IPsec engines NP6 processors use	1272
Stripping clear text padding and IPsec session ESP padding	1273
Disable NP6 CAPWAP offloading	1273
Optionally disable NP6 offloading of traffic passing between 10Gbps and 1Gbps interfaces	1273
Offloading RDP traffic	1274
NP6 session drift	1274
Optimizing FortiGate-3960E and 3980E IPsec VPN performance	1275
Recalculating packet checksums if the iph.reserved bit is set to 0	1275
FortiGate NP6 architectures	1276
FortiGate-300D fast path architecture	1276
FortiGate-300E and 301E fast path architecture	1276
FortiGate-400D fast path architecture	1278
FortiGate-500D fast path architecture	1279
FortiGate-500E and 501E fast path architecture	1279
FortiGate-600D fast path architecture	1280
FortiGate-800D fast path architecture	1281
FortiGate-900D fast path architecture	1283
FortiGate-1000D fast path architecture	1285
FortiGate-1200D fast path architecture	1287
FortiGate-1500D fast path architecture	1290
FortiGate-1500DT fast path architecture	1291
FortiGate-2000E fast path architecture	1293
FortiGate-2500E fast path architecture	1295
FortiGate-3000D fast path architecture	1298
FortiGate-3100D fast path architecture	1299
FortiGate-3200D fast path architecture	1300
FortiGate-3700D fast path architecture	1303
FortiGate-3700DX fast path architecture	1307
FortiGate-3800D fast path architecture	1311
FortiGate-3810D fast path architecture	1313
FortiGate-3815D fast path architecture	1314
FortiGate-3960E fast path architecture	1316
FortiGate-3980E fast path architecture	1317
FortiGate-5001D fast path architecture	1319
FortiGate-5001E and 5001E1 fast path architecture	1321
FortiController-5902D fast path architecture	1323
FortiGate NP6lite architectures	1326
FortiGate-200E and 201E fast path architecture	1326
NP4 and NP4Lite acceleration	1328

Viewing your FortiGate NP4 processor configuration.....	1328
NP4 and NP4Lite processors and sFlow and NetFlow.....	1329
Configuring NP4 traffic offloading.....	1329
Disabling NP4 and NP4lite hardware acceleration (fastpath).....	1331
Increasing NP4 offloading capacity using link aggregation groups (LAGs).....	1331
NP4 traffic shaping offloading.....	1331
NP4 IPsec VPN offloading.....	1332
Configuring inter-VDOM link acceleration with NP4 processors.....	1332
Offloading NP4 anomaly detection.....	1336
FortiGate NP4 architectures.....	1339
FortiGate-600C.....	1339
FortiGate-800C.....	1340
FortiGate-1000C.....	1341
FortiGate-1240B.....	1342
FortiGate-3040B.....	1343
FortiGate-3140B.....	1344
FortiGate-3140B — load balance mode.....	1345
FortiGate-3240C.....	1345
FortiGate-3600C.....	1347
FortiGate-3950B and FortiGate-3951B.....	1348
FortiGate-3950B and FortiGate-3951B — load balance mode.....	1350
FortiGate-5001C.....	1351
FortiGate-5001B.....	1352
Setting switch-mode mapping on the ADM-XD4.....	1353
Hardware acceleration get and diagnose commands.....	1354
get hardware npu np6.....	1354
diagnose npu np6.....	1354
Using diagnose npu np6 npu-feature to verify enabled NP6 features.....	1355
Using the diagnose sys session/session6 list command.....	1356
diagnose npu np6 session-stats <np6-id> (number of NP6 IPv4 and IPv6 sessions).....	1359
diagnose npu np6 ipsec-stats (NP6 IPsec statistics).....	1361
diagnose sys mcast-session/session6 list (IPv4 and IPv6 multicast sessions).....	1362
diagnose npu np6 sse-stats <np6-id> (number of NP6 sessions and dropped sessions).....	1363
diagnose npu np6 dce <np6-id> (number of dropped NP6 packets).....	1363
diagnose hardware deviceinfo nic <interface-name> (number of packets dropped by an interface).....	1364
diagnose npu np6 synproxy-stats (NP6 SYN-proxied sessions and unacknowledged SYNs).....	1365
Chapter 12 - High Availability.....	1366
What's new in FortiOS 6.0.....	1366
Solving the high availability problem.....	1367

FortiGate Cluster Protocol (FGCP).....	1367
FortiGate Session Life Support Protocol (FGSP).....	1368
VRRP high availability.....	1369
Session-Aware Load Balancing Clustering (SLBC).....	1370
Enhanced Load Balancing Clustering (ELBC).....	1371
Content clustering.....	1371
An introduction to the FGCP.....	1373
About the FGCP.....	1373
Synchronizing the configuration (and settings that are not synchronized).....	1376
Preparing the FortiGates before setting up an FGCP cluster.....	1377
Configuring FortiGates for FGCP HA operation.....	1378
Active-passive and active-active HA.....	1382
Identifying the cluster and cluster units.....	1383
Device failover, link failover, and session failover.....	1384
Primary unit selection.....	1385
HA override.....	1393
FortiGate HA compatibility with DHCP and PPPoE.....	1397
HA and distributed clustering.....	1398
Clusters of three or four FortiGates.....	1399
Disk storage configuration and HA.....	1402
FGCP high availability best practices.....	1403
FGCP HA terminology.....	1405
FGCP support for OCVPN.....	1408
HA GUI options.....	1409
FGCP configuration examples and troubleshooting.....	1412
About the examples in this chapter.....	1412
How to set up FGCP clustering (recommended steps).....	1412
Adding a third FortiGate to an operating cluster and switching to active-active HA.....	1412
Active-active HA cluster in transparent mode.....	1412
FortiGate-5000 active-active HA cluster with FortiClient licenses.....	1424
Replacing a failed cluster unit.....	1434
FGCP HA with 802.3ad aggregated interfaces.....	1436
Example HA and redundant interfaces.....	1447
Troubleshooting HA clusters.....	1458
Virtual clusters.....	1463
Virtual clustering overview.....	1463
Configuring virtual clustering.....	1465
Virtual clustering configuration examples.....	1469
Inter-VDOM links in a virtual clustering configuration.....	1470
Troubleshooting virtual clustering.....	1472
Full mesh HA.....	1473
Full mesh HA overview.....	1473

Example full mesh HA configuration.....	1475
Troubleshooting full mesh HA.....	1485
Operating clusters and virtual clusters.....	1487
Operating a cluster.....	1487
Operating a virtual cluster.....	1487
Managing individual cluster units using a reserved out-of-band management interface.....	1488
Managing individual cluster units using an in-band management IP address.....	1494
Managing individual cluster units in a virtual cluster.....	1494
Shutting down or rebooting the primary unit.....	1495
The primary unit acts as a router for subordinate unit management traffic.....	1495
Cluster communication with RADIUS and LDAP servers.....	1496
Clusters and FortiGuard services.....	1496
Clusters and logging.....	1497
Clusters and SNMP.....	1500
Adding FortiClient licenses to a cluster.....	1503
Cluster members list.....	1505
Virtual cluster members list.....	1505
Viewing HA statistics.....	1506
Changing the HA configuration of an operating cluster.....	1506
Changing the HA configuration of an operating virtual cluster.....	1506
Changing the subordinate unit host name and device priority.....	1507
Upgrading cluster firmware.....	1507
Downgrading cluster firmware.....	1509
Backing up and restoring the cluster configuration.....	1510
Monitoring cluster units for failover.....	1510
Viewing cluster status from the CLI.....	1511
Managing individual cluster units.....	1518
Disconnecting a cluster unit from a cluster.....	1518
Adding a disconnected FortiGate back to its cluster.....	1519
diagnose sys ha dump-by command.....	1520
HA and failover protection.....	1521
About active-passive failover.....	1521
About active-active failover.....	1522
Device failover.....	1522
HA heartbeat and communication between cluster units.....	1523
Cluster virtual MAC addresses.....	1530
Synchronizing the configuration.....	1537
Synchronizing kernel routing tables.....	1546
Configuring graceful restart for dynamic routing failover.....	1547
Link failover (port monitoring or interface monitoring).....	1549
Monitoring VLAN interfaces.....	1555

Sub-second failover.....	1555
Remote link failover.....	1556
Failover and attached network equipment.....	1561
Monitoring cluster units for failover.....	1561
NAT/Route mode active-passive cluster packet flow.....	1561
Transparent mode active-passive cluster packet flow.....	1563
Failover performance.....	1566
Session failover (session-pickup).....	1568
Enabling session-pickup for TCP, UDP, ICMP, and multicast session failover.....	1568
If session pickup is disabled.....	1569
Improving session synchronization performance.....	1569
Session failover limitations for sessions passing through the cluster.....	1570
Session failover limitations for sessions terminated by the cluster.....	1573
Synchronizing IPsec VPN SAs.....	1575
WAN optimization and HA.....	1576
HA and load balancing.....	1577
Load balancing overview.....	1577
Selecting a load balancing schedule.....	1579
Load balancing TCP and UDP sessions.....	1580
Using NP4 or NP6 processors to offload load balancing.....	1580
Configuring weighted-round-robin weights.....	1580
Dynamically optimizing weighted load balancing according to how busy cluster units are.....	1581
Example weighted load balancing configuration.....	1583
NAT/Route mode active-active cluster packet flow.....	1585
Transparent mode active-active cluster packet flow.....	1588
HA with FortiGate-VM and third-party products.....	1592
FortiGate-VM for VMware HA configuration.....	1592
FortiGate-VM for Hyper-V HA configuration.....	1592
Troubleshooting layer-2 switches.....	1593
Failover issues with layer-3 switches.....	1593
Failover and attached network equipment.....	1593
Ethertype conflicts with third-party switches.....	1594
LACP, 802.3ad aggregation and third-party switches.....	1594
VRRP high availability.....	1595
Configuring VRRP.....	1596
Adding an IPv4 VRRP virtual router to a FortiGate interface.....	1596
Adding an IPv6 VRRP virtual router to a FortiGate interface.....	1597
Setting up VRRP failover.....	1597
Using VRRP virtual MAC addresses.....	1599
Setting up VRRP groups.....	1600
Example IPv4 VRRP configuration: two FortiGates in a VRRP group.....	1601

Example IPv4 VRRP configuration: VRRP load balancing two FortiGates and two VRRP groups.....	1603
Optional VRRP configuration settings.....	1604
FortiController-5000 SLBC support.....	1604
FortiGate Session Life Support Protocol (FGSP).....	1606
Session synchronization between FGCP clusters.....	1608
Configuring FGSP HA cluster-sync instances.....	1609
Synchronizing TCP and SCTP sessions.....	1611
Synchronizing the configuration.....	1611
FGSP and firmware upgrades.....	1612
Backing up and restoring the configuration of an FGSP cluster.....	1613
IPsec tunnel synchronization.....	1613
Synchronizing UDP and ICMP (connectionless) sessions.....	1613
Synchronizing NAT sessions.....	1613
Synchronizing asymmetric sessions.....	1614
Synchronizing expectation sessions.....	1615
GTP session synchronization: FGSP for FortiOS Carrier.....	1615
Security profile flow-based inspection and asymmetric traffic.....	1615
Notes and limitations.....	1616
Configuring session synchronization links.....	1616
Basic example configuration.....	1617
Verifying the FGSP configuration and synchronization.....	1619
Chapter 13 - IPsec VPN.....	1621
Introduction.....	1621
What's new in FortiOS 6.0.2.....	1622
What's new in FortiOS 6.0.1.....	1622
What's new in FortiOS 6.0.....	1622
IPsec VPN concepts.....	1624
VPN tunnels.....	1624
VPN gateways.....	1626
Clients, servers, and peers.....	1628
Encryption.....	1629
Authentication.....	1631
Phase 1 and Phase 2 settings.....	1632
Security Association.....	1633
IKE and IPsec packet processing.....	1633
IPsec VPN overview.....	1638
Types of VPNs.....	1638
Planning your VPN.....	1639
General preparation steps.....	1641
How to use this guide to configure an IPsec VPN.....	1641
IPsec VPN in the web-based manager.....	1642

Phase 1 configuration.....	1642
Phase 2 configuration.....	1649
Concentrator.....	1653
IPsec Monitor.....	1654
Phase 1 parameters.....	1655
Overview.....	1655
Defining the tunnel ends.....	1655
Choosing Main mode or Aggressive mode.....	1656
Choosing the IKE version.....	1656
Authenticating the FortiGate unit.....	1658
Authenticating remote peers and clients.....	1660
Defining IKE negotiation parameters.....	1666
Using XAuth authentication.....	1672
Dynamic IPsec route control.....	1673
Phase 2 parameters.....	1675
Phase 2 settings.....	1675
Configuring the Phase 2 parameters.....	1677
Defining VPN security policies.....	1681
Defining policy addresses.....	1681
Defining security policies for policy-based and route-based VPNs.....	1682
Gateway-to-gateway.....	1687
Configuration overview.....	1687
Gateway-to-gateway configuration.....	1689
How to work with overlapping subnets.....	1694
Testing.....	1699
Hub-and-spoke configurations.....	1702
Configuration overview.....	1702
Configure the hub.....	1704
Configure the spokes.....	1709
Dynamic spokes configuration example.....	1712
One-Click VPN (OCVPN).....	1719
General configuration.....	1719
Key exchange.....	1720
Device polling and controller information.....	1720
System states.....	1721
Debugging and logging.....	1721
Dynamic DNS configuration.....	1723
Dynamic DNS over VPN concepts.....	1723
DDNS topology.....	1725
Configuration overview.....	1726
FortiClient dialup-client configuration.....	1736
Configuration overview.....	1736

FortiGate dialup-client configurations	1745
Configuration overview	1745
Supporting IKE Mode Config clients	1753
IKE Mode Config overview	1753
Automatic configuration overview	1753
IKE Mode Config method	1753
Internet-browsing configuration	1759
Configuration overview	1759
Routing all remote traffic through the VPN tunnel	1761
Redundant VPN configurations	1763
Configuration overview	1763
IPsec VPN tunnel aggregate interfaces	1767
Transparent mode VPNs	1769
Configuration overview	1769
IPv6 IPsec VPNs	1774
Configuration examples	1775
L2TP and IPsec (Microsoft VPN)	1786
Overview	1786
Assumptions	1787
Configuration overview	1787
GRE over IPsec (Cisco VPN)	1795
Configuration overview	1796
Configuring the Cisco router	1801
Keep-alive support for GRE	1802
Protecting OSPF with IPsec	1803
Configuration overview	1803
Redundant OSPF routing over IPsec	1810
OSPF over dynamic IPsec	1814
BGP over dynamic IPsec	1817
IPsec Auto-Discovery VPN (ADVPN)	1821
Example ADVPN configuration	1822
Logging and monitoring	1827
Monitoring VPN connections	1827
VPN event logs	1828
Troubleshooting	1829
Common IPsec VPN problems	1829
Troubleshooting connection issues	1833
General troubleshooting tips	1837
Troubleshooting L2TP and IPsec	1838
Troubleshooting GRE over IPsec	1841
Chapter 14 - Logging and Reporting	1844
What's new in FortiOS 6.0	1845

Automatic synchronization of log display location.....	1845
Improved log messages for SD-WAN link quality changes.....	1845
Extended UTM logging and improved syslog configuration.....	1845
Updated reliable syslog encryption to comply with RFC 5425.....	1845
Improved log display consistency at high load.....	1846
Logging and reporting overview.....	1847
What is logging?.....	1847
FortiOS features available for logging.....	1848
Log messages.....	1853
Log files and types.....	1859
Log database and datasets.....	1860
Log devices.....	1862
Reports.....	1866
Best practices: Log management.....	1867
Logging and reporting for small networks.....	1869
Modifying default log device settings.....	1869
Configuring the backup solution.....	1871
Logging and reporting for large networks.....	1873
Modifying default log device settings.....	1873
Configuring the backup solution.....	1875
Advanced logging.....	1878
Log backup and restore tools.....	1878
Configuring logging to multiple Syslog servers.....	1878
Using Automatic Discovery to connect to a FortiAnalyzer unit.....	1879
Activating a FortiCloud account for logging purposes.....	1880
Viewing log storage space.....	1880
Customizing and filtering log messages.....	1881
Viewing logs from the CLI.....	1881
Configuring NAC Quarantine logging.....	1882
Logging local-in policies.....	1882
Tracking specific search phrases in reports.....	1884
Interpreting and configuring FSSO syslog log messages.....	1885
Troubleshooting and logging.....	1887
Using log messages to help in troubleshooting issues.....	1887
Connection issues between FortiGate unit and logging devices.....	1888
Log database issues.....	1888
Logging daemon (Miglogd).....	1890
Chapter 15 - Managing Devices.....	1891
What's new in FortiOS 6.0.....	1891
Managing “bring your own device”.....	1892
Device monitoring.....	1892
Device groups.....	1894

Controlling access with a MAC Address Access Control List	1894
Security policies for devices	1896
Chapter 16 - FortiSwitch Devices Managed by FortiOS 6.0.1	1900
Introduction	1900
Supported models	1900
Support of FortiLink features	1902
Before you begin	1903
How this guide is organized	1903
What's new in FortiOS 6.0.1	1904
What's new in FortiOS 6.0	1904
FortiOS 6.0.1	1904
FortiOS 6.0	1907
Connecting FortiLink ports	1919
1. Enable the switch controller on the FortiGate unit	1919
2. Connect the FortiSwitch unit and FortiGate unit	1919
FortiLink configuration using the FortiGate GUI	1921
Summary of the procedure	1921
FortiLink split interface	1922
Authorizing the FortiSwitch unit	1922
Adding preauthorized FortiSwitch units	1922
Managed FortiSwitch display	1923
Edit a managed FortiSwitch unit	1923
Network interface display	1924
Add link aggregation groups (Trunks)	1924
Configure DHCP blocking, IGMP snooping, STP, and loop guard on managed FortiSwitch ports	1924
FortiLink configuration using the FortiGate CLI	1926
Summary of the procedure	1926
Configure FortiLink on a physical port	1926
Configure FortiLink on a logical interface	1927
Enable multiple FortiLink interfaces	1928
FortiLink mode over a layer-3 network	1928
Network topologies for managed FortiSwitch units	1930
Supported topologies	1930
Single FortiGate managing a single FortiSwitch unit	1931
Single FortiGate unit managing a stack of several FortiSwitch units	1932
HA-mode FortiGate units managing a single FortiSwitch unit	1933
HA-mode FortiGate units managing a stack of several FortiSwitch units	1934
HA-mode FortiGate units managing a FortiSwitch two-tier topology	1935
Single FortiGate unit managing multiple FortiSwitch units (using a hardware or software switch interface)	1936
HA-mode FortiGate units managing two-tier FortiSwitch units with access rings	1937

Dual-homed servers connected to FortiLink tier-1 FortiSwitch units using an MCLAG	1938
Standalone FortiGate unit with dual-homed FortiSwitch access	1939
HA-mode FortiGate units with dual-homed FortiSwitch access	1940
Multi-tiered MCLAG with HA-mode FortiGate units	1941
Grouping FortiSwitch units	1946
Transitioning from a FortiLink split interface to a FortiLink MCLAG	1950
Optional setup tasks	1952
Configuring the FortiSwitch management port	1952
Converting to FortiSwitch standalone mode	1953
Changing the admin password on the FortiGate for all managed FortiSwitch units	1953
Enabling network-assisted device detection	1954
Limiting the number of parallel process for FortiSwitch configuration	1954
FortiSwitch features configuration	1954
Configure VLANs	1955
Configure IGMP settings	1958
Configure LLDP-MED	1958
Configure the MAC sync interval	1960
Configure STP settings	1960
Quarantines	1961
FortiSwitch port features	1966
FortiSwitch ports display	1966
Configuring ports using the GUI	1967
Configuring ports using the FortiGate CLI	1967
FortiSwitch port security policy	1981
Configure the 802.1X settings for a virtual domain	1982
Override the virtual domain settings	1983
Define an 802.1X security policy	1983
Apply an 802.1X security policy to a FortiSwitch port	1985
Test 802.1x authentication with monitor mode	1985
Restrict the type of frames allowed through IEEE 802.1Q ports	1986
RADIUS accounting support	1986
Additional capabilities	1987
Execute custom FortiSwitch commands	1987
View and upgrade the FortiSwitch firmware version	1988
FortiSwitch log export	1989
FortiSwitch per-port device visibility	1989
FortiGate CLI support for FortiSwitch features (on non-FortiLink ports)	1989
Synchronizing the FortiGate unit with the managed FortiSwitch units	1994
Replacing a managed FortiSwitch unit	1995
Troubleshooting	1996
Troubleshooting FortiLink issues	1996
Chapter 17 - Networking	1998

Introduction.....	1998
What's new in FortiOS 6.0.....	1998
Interfaces.....	1999
Administrative access.....	1999
Aggregate interfaces.....	2000
DHCP addressing mode on an interface.....	2001
DHCP servers and relays.....	2002
Interface MTU packet size.....	2010
Interface settings.....	2011
Loopback interfaces.....	2016
One-armed sniffer.....	2017
Physical ports.....	2020
PPPoE addressing mode on an interface.....	2021
Probing interfaces.....	2023
Redundant interfaces.....	2024
Dual Internet connections.....	2025
Secondary IP addresses to an interface.....	2028
Software switch.....	2028
Soft switch example.....	2029
Virtual switch.....	2030
Zones.....	2033
Virtual domains.....	2034
Wireless.....	2035
VLANs.....	2035
Enhanced MAC VLANs.....	2049
Virtual wire pairs.....	2052
Botnet and command-and-control protection.....	2053
DNS.....	2055
DNS settings.....	2055
Additional DNS CLI configuration.....	2055
DDNS.....	2056
DNS servers.....	2057
Advanced static routing.....	2062
Routing concepts.....	2062
Static routing tips.....	2078
Policy routing.....	2078
Static routing in transparent mode.....	2083
Static routing example.....	2084
Dynamic routing.....	2093
Overview.....	2093
Comparison of dynamic routing protocols.....	2096
Choosing a routing protocol.....	2101

Dynamic routing terminology.....	2102
Controlling how routing changes affect active sessions.....	2108
IPv6 in dynamic routing.....	2109
RIP.....	2110
OSPF.....	2146
BGP.....	2187
IS-IS.....	2224
Multicast forwarding.....	2241
Sparse mode.....	2241
Dense mode.....	2242
PIM support.....	2243
Multicast forwarding and FortiGate devices.....	2244
Configuring FortiGate multicast forwarding.....	2245
Multicast routing examples.....	2248
SD-WAN.....	2273
Application control settings in SD-WAN rules.....	2273
Internet services in SD-WAN rules.....	2273
Bandwidth and custom profile options in SD-WAN rules.....	2274
SLA management.....	2275
Multiple server support for health checks.....	2276
IPv6 support for SD-WAN.....	2277
DSCP tagging of forwarded packets in SD-WAN rules.....	2278
SD-WAN and dynamic routing.....	2278
Link priority in SD-WAN rules.....	2279
SD-WAN rule for address negation.....	2279
SD-WAN CLI changes.....	2280
Troubleshooting.....	2281
Netflow support.....	2281
sFlow support.....	2281
Packet capture.....	2283
Chapter 18 - Parallel Path Processing (Life of a Packet).....	2284
Parallel Path Processing.....	2285
High-level list of processes that affect packets.....	2286
Packet flow ingress and egress: FortiGates without network processor offloading.....	2287
Ingress.....	2287
Admission control.....	2288
Kernel.....	2288
UTM/NGFW.....	2290
Content processors (CP8 and CP9).....	2290
Kernel.....	2292
Egress.....	2292
Packet flow: FortiGates with NP6 processors first packet of a new session.....	2293

Network processors (NP6).....	2294
Packet flow: FortiGates with NP6 processors - packets in an offloaded session.....	2295
Packet flow: FortiGates with NP6 processors - packets in an NTurbo session.....	2296
UTM/NGFW packet flow: flow-based inspection.....	2298
UTM/NGFW packet flow: proxy-based inspection.....	2300
UTM/NGFW packet flow: explicit web proxy.....	2302
Comparison of inspection types.....	2304
Mapping security functions to inspection types.....	2304
More information about inspection methods.....	2304
Chapter 19 - Sandbox Inspection.....	2306
An Overview of Sandbox Inspection.....	2307
What is Sandbox Inspection?.....	2307
FortiSandbox Appliance vs FortiSandbox Cloud.....	2307
Sending Files for Sandbox Inspection.....	2308
Using FortiSandbox with a FortiGate.....	2309
Connecting a FortiGate to FortiSandbox.....	2309
FortiSandbox Console.....	2310
Sandbox Integration.....	2311
Overview.....	2311
Example Configuration.....	2312
Sandbox Inspection FAQ.....	2315
Chapter 20 - Fortinet Security Fabric.....	2316
Introduction.....	2316
What's new in FortiOS 6.0.2.....	2316
What's new in FortiOS 6.0.1.....	2316
What's new in FortiOS 6.0.0.....	2317
Fortinet Security Fabric overview.....	2318
Access security.....	2319
Client security.....	2319
Application security.....	2319
Cloud security.....	2319
NOC and SOC security.....	2320
Advanced threat intelligence.....	2320
Partner API.....	2320
The Security Fabric solution components.....	2321
Devices in the Security Fabric.....	2321
Security Fabric topology views.....	2324
Security Fabric Rating.....	2325
FortiTelemetry.....	2325
Configuring the Fortinet Security Fabric.....	2326
Forming the Security Fabric.....	2326
Setting up data collection with FortiAnalyzer.....	2331

Adding a FortiSandbox to the Security Fabric.....	2333
Adding a FortiManager to the Security Fabric.....	2334
Adding FortiClient EMS to the Security Fabric.....	2335
Using the Fortinet Security Fabric.....	2337
Understanding the Security Fabric dashboard widgets.....	2337
Viewing the Security Fabric topology.....	2339
Running a Security Fabric Rating.....	2344
Automation stitches.....	2349
Trigger events.....	2349
Response actions.....	2351
Creating automation stitches.....	2353
Configuring an automation, trigger, and action in the CLI.....	2355
Chaining and delaying actions for AWS Lambda and webhook.....	2359
Diagnose commands for automation stitches.....	2359
Fabric Connectors.....	2362
Available services for Fabric Connectors.....	2362
Configuring Fabric Connectors.....	2363
Verifying Fabric Connector status.....	2366
Central management with FortiManager.....	2368
Configuring the FortiManager.....	2368
FortiGuard.....	2370
Firmware updates.....	2377
Administrative domains.....	2377
Backing up and restoring configurations.....	2377
FortiManager in backup mode.....	2377
Related resources.....	2379
Chapter 21 - Security Profiles.....	2380
What's new in FortiOS 6.0.1.....	2381
What's new in FortiOS 6.0.....	2381
Inside FortiOS: AntiVirus.....	2382
Advanced protection against malware and APTs.....	2382
Features.....	2382
Inside FortiOS: Application Control.....	2386
Enhance control and network visibility.....	2386
Features.....	2387
Recipes.....	2389
Inside FortiOS: Intrusion Prevention System (IPS).....	2390
World class next generation IPS capabilities.....	2390
Features.....	2390
Inside FortiOS: Web Filtering.....	2395
Intelligent and effective content control.....	2395
Features.....	2396

Security profiles overview.....	2400
Traffic inspection.....	2400
Content inspection and filtering.....	2402
Security profile components.....	2405
Security profiles/lists/sensors.....	2407
Inspection modes.....	2408
Proxy-based inspection.....	2408
Flow-based inspection.....	2408
Changing between proxy and flow mode.....	2409
Comparison of inspection types.....	2410
Individual security profile considerations.....	2411
Proxy mode and flow mode antivirus and web filter profile options.....	2411
AntiVirus.....	2416
Antivirus concepts.....	2416
Enabling AntiVirus scanning.....	2428
Testing your antivirus configuration.....	2434
Example scenarios.....	2434
Web filter.....	2441
Web filter concepts.....	2441
Inspection modes.....	2443
FortiGuard Web Filtering Service.....	2444
Configuring web filter profiles.....	2450
Overriding FortiGuard website categorization.....	2455
Using cookies to authenticate users in a Web Filter override.....	2462
Web Profile Overrides.....	2463
SafeSearch.....	2464
YouTube Education Filter.....	2465
Static URL filter.....	2466
Web content filter.....	2471
Web filtering example.....	2474
Advanced web filter configurations.....	2477
DNS filter.....	2482
Application control.....	2485
Application control concepts.....	2485
Enabling application control in profile-based modes.....	2486
Application control actions.....	2489
Application considerations.....	2490
Application control monitor.....	2490
Application control examples.....	2491
Intrusion prevention.....	2495
IPS concepts.....	2495
Enabling IPS scanning.....	2498

IPS processing in an HA cluster.....	2501
Configure IPS options.....	2501
Enabling IPS packet logging.....	2506
Other IPS examples.....	2506
Anti-spam filter.....	2513
Anti-spam concepts.....	2513
Anti-spam techniques.....	2513
Configuring Anti-spam.....	2517
Order of spam filtering.....	2519
Spam actions.....	2520
Anti-spam examples.....	2521
Data leak prevention.....	2523
Data leak prevention concepts.....	2523
Enable data leak prevention.....	2528
Creating/editing a DLP sensor.....	2529
DLP archiving.....	2531
DLP examples.....	2532
ICAP support.....	2537
The protocol.....	2537
Offloading using ICAP.....	2538
Configuring ICAP.....	2538
Example ICAP sequence.....	2539
Example ICAP scenario.....	2539
FortiClient Compliance Profiles.....	2541
Endpoint protection overview.....	2541
Configuring endpoint protection.....	2543
Configuring endpoint registration over a VPN.....	2548
Assigning FortiClient Profiles using Microsoft AD user groups.....	2550
Modifying the endpoint protection replacement messages.....	2552
Monitoring endpoints.....	2552
Proxy options.....	2554
The use of different proxy profiles and profile options.....	2554
Proxy Options profile components.....	2554
SSL/SSH inspection.....	2557
SSL inspection.....	2557
Why use SSL inspection.....	2560
SSL certificate inspection.....	2561
Creating or editing an SSL/SSH Inspection profile.....	2562
Secure white list database.....	2564
SSH MITM deep inspection.....	2564
SSL server table for SSL offloading.....	2568
Custom Application & IPS Signatures.....	2569

Creating a custom IPS signature.....	2569
Custom signature syntax.....	2569
Custom signature keywords.....	2570
Creating a custom signature to block access to example.com.....	2582
Creating a custom signature to block the SMTP “vrfy” command.....	2584
Creating a custom signature to block files according to the file's hash value.....	2585
Other security profiles considerations.....	2587
Global security profiles across Virtual domains (VDOMs).....	2587
Conserve mode.....	2587
Using wildcards and Perl regular expressions.....	2589
Control how sessions are distributed to Fortinet processes.....	2592
CPU allocation and tuning commands to survive reboot.....	2592
Excluding industrial IP signatures.....	2592
Chapter 22 - Server Load Balancing.....	2594
Inside FortiOS: Server Load Balancing.....	2594
Server Load Balancing combined with NGFW and UTM protection.....	2594
SSL/TLS offloading.....	2595
SSL/TLS content inspection.....	2596
Health Check.....	2596
Server Monitoring and Management.....	2596
HTTP Multiplexing.....	2596
Basic load balancing configuration example.....	2597
Configuring load balancing.....	2601
Load balancing and other FortiOS features.....	2603
Configuring load balancing from the GUI.....	2603
Configuring load balancing from the CLI.....	2605
Load balancing methods.....	2606
Session persistence.....	2607
Real servers.....	2607
Health check monitoring.....	2609
Load balancing limitations.....	2611
Monitoring load balancing.....	2612
Load balancing diagnose commands.....	2613
HTTP and HTTPS load balancing, multiplexing, and persistence.....	2614
HTTP and HTTPS multiplexing.....	2614
HTTP and HTTPS persistence.....	2615
HTTP host-based load balancing.....	2617
SSL/TLS load balancing.....	2619
SSL/TLS offloading.....	2619
Separate virtual-server client and server TLS version and cipher configuration.....	2621
Setting the SSL/TLS versions to use for server and client connections.....	2621
Setting the SSL/TLS cipher choices for server and client connections.....	2622

Protection from TLS protocol downgrade attacks.....	2623
Setting 3072- and 4096-bit Diffie-Hellman values.....	2623
Additional SSL load balancing and SSL offloading options.....	2623
SSL offloading support for Internet Explorer 6.....	2624
Selecting the cipher suites available for SSL load balancing.....	2625
Disabling SSL/TLS re-negotiation.....	2625
IP, TCP, and UDP load balancing.....	2629
Example HTTP load balancing to three real web servers.....	2629
GUI configuration.....	2630
CLI configuration.....	2633
Example Basic IP load balancing configuration.....	2634
Example Adding a server load balance port forwarding virtual IP.....	2635
Example Weighted load balancing configuration.....	2636
GUI configuration.....	2636
CLI configuration.....	2638
Example HTTP and HTTPS persistence configuration.....	2639
CLI configuration: adding persistence for a specific domain.....	2642
Chapter 23 - SSL VPN.....	2644
What's new in FortiOS 6.0.1.....	2644
What's new in FortiOS 6.0.....	2644
Overview.....	2645
SSL VPN modes of operation.....	2646
Port forwarding mode.....	2647
SSL VPN conserve mode.....	2649
Traveling and security.....	2650
SSL VPN and IPv6.....	2650
Basic configuration.....	2651
User accounts and groups.....	2651
Configuring SSL VPN web portals.....	2656
Configuring security policies.....	2665
Configuring encryption key algorithms.....	2671
Additional configuration options.....	2672
The SSL VPN client.....	2683
FortiClient.....	2683
Tunnel mode client configuration.....	2683
The SSL VPN web portal.....	2685
Connecting to the FortiGate unit.....	2685
Web portal overview.....	2685
Portal configuration.....	2687
Using the Bookmarks widget.....	2692
Using the Quick Connection Tool.....	2694
Using FortiClient.....	2698

Setup examples	2699
Secure Internet browsing	2699
Split tunnel	2701
Multiple user groups with different access permissions	2704
Client device certificate authentication with multiple groups	2709
Troubleshooting	2710
Chapter 24 - System Administration	2713
What's new in FortiOS 6.0	2713
Administrators	2714
Administrator profiles	2714
Adding a local administrator	2715
LDAP authentication for administrators	2716
Other methods of administrator authentication	2717
Administrator lockout	2718
Monitoring administrators	2719
Management access	2719
Security precautions	2720
Monitoring	2723
Dashboard	2723
sFlow support	2723
Monitor menus	2725
Logging	2725
Alert email	2726
SNMP	2727
SNMP get command syntax	2735
Replacement messages	2737
Administration for schools	2743
PPTP and L2TP	2746
How PPTP VPNs work	2746
FortiGate unit as a PPTP server	2748
Configuring the FortiGate unit for PPTP VPN	2751
Configuring the FortiGate unit for PPTP passthrough	2751
Testing PPTP VPN connections	2753
Logging VPN events	2753
Configuring L2TP VPNs	2754
L2TP configuration overview	2756
Session helpers	2760
Viewing the session helper configuration	2760
Changing the session helper configuration	2761
DCE-RPC session helper (dcerpc)	2763
DNS session helpers (dns-tcp and dns-udp)	2764
File transfer protocol (FTP) session helper (ftp)	2764

H.323 and RAS session helpers (h323 and ras).....	2764
H.245 session helper (h245).....	2765
Media Gateway Controller Protocol (MGCP) session helper (mgcp).....	2765
ONC-RPC portmapper session helper (pmap).....	2766
PPTP session helper for PPTP traffic (pptp).....	2766
Remote shell session helper (rsh).....	2767
Real-Time Streaming Protocol (RTSP) session helper (rtsp).....	2767
Session Initiation Protocol (SIP) session helper (sip).....	2768
Trivial File Transfer Protocol (TFTP) session helper (tftp).....	2768
Oracle TNS listener session helper (tns).....	2768
Advanced concepts.....	2769
Single firewall vs. multiple virtual domains.....	2769
Modem.....	2771
Assigning IP address by MAC address.....	2774
IP addresses for self-originated traffic.....	2775
Disk.....	2775
CLI scripts.....	2776
Rejecting PING requests.....	2778
Opening TCP 113.....	2778
Obfuscate HTTP responses from SSL VPN.....	2779
Blocking land attacks in transparent mode.....	2779
Multi-dimension tagging.....	2779
Chapter 25 - Traffic Shaping.....	2781
What's new in FortiOS 6.0.....	2781
The purpose of traffic shaping.....	2782
QoS.....	2782
Traffic policing.....	2783
Bandwidth guarantee, limit, and priority interactions.....	2784
FortiGate traffic.....	2785
Through traffic.....	2785
Important considerations.....	2789
Traffic shaping methods.....	2791
Traffic shaping options.....	2791
Shared policy traffic shaping.....	2792
Per-IP traffic shaping.....	2797
Application control shaping.....	2800
Reverse direction traffic shaping.....	2802
Enabling traffic shaping in the security policy.....	2803
Scheduling traffic shaping policies.....	2803
ToS priority.....	2804
Interface-based traffic shaping.....	2805
Differentiated services.....	2811

DSCP examples.....	2812
Traffic mapping.....	2816
Traffic shaper monitor.....	2817
Examples.....	2818
QoS using priority from security policies.....	2818
QoS using priority from ToS or DiffServ.....	2822
Example setup for VoIP.....	2823
Troubleshooting traffic shaping.....	2829
Interface diagnosis.....	2829
Traffic shaper diagnose commands.....	2829
Packet lost with the debug flow.....	2831
Session list details with dual traffic shaper.....	2831
Additional information.....	2831
Chapter 26 - Transparent Mode.....	2833
What's new in FortiOS 6.0.....	2833
Transparent mode overview.....	2834
What is transparent mode?.....	2834
Transparent mode features.....	2834
Installation.....	2836
Installing a FortiGate in transparent mode.....	2836
Using a virtual wire pair to simplify transparent mode.....	2837
Management IP configuration.....	2838
Networking in transparent mode.....	2840
Static routing.....	2840
Packet forwarding.....	2841
Network address translation (NAT).....	2846
VLANs and forwarding domains.....	2848
Inter-VDOM links between NAT/route and transparent VDOMs.....	2851
Replay traffic scenario.....	2851
Packet forwarding using Cisco protocols.....	2852
Configuration example.....	2853
Firewalls and security in transparent mode.....	2856
Firewall policy look up.....	2856
Firewall session list.....	2856
Security scanning.....	2857
IPsec VPN in transparent mode.....	2858
Using IPsec VPNs in transparent mode.....	2858
Example 1: Remote sites with different subnets.....	2859
Example 2: Remote sites on the same subnet.....	2865
Using FortiManager and FortiAnalyzer.....	2872
High availability in transparent mode.....	2873
Virtual clustering.....	2873

MAC address assignment	2874
Best practices	2875
Chapter 27 - Troubleshooting	2876
Troubleshooting methodologies	2877
Ensure you have administrator-level access to required equipment	2877
Establish a baseline	2877
Define the problem	2878
Create a troubleshooting plan	2879
Obtain any required equipment	2880
Consult Fortinet resources	2880
Troubleshooting tools	2881
FortiOS diagnostics	2881
FortiOS ports	2903
FortiAnalyzer and FortiManager ports	2904
FortiGuard troubleshooting	2904
Troubleshooting tips	2908
How to check hardware connections	2909
How to check FortiOS network settings	2910
How to check CPU and memory resources	2911
How to check modem status	2917
How to run ping and traceroute	2917
How to check the logs	2921
How to verify the contents of the routing table (in NAT mode)	2922
How to verify the correct route is being used	2923
How to verify the correct firewall policy is being used	2924
How to check the bridging information in transparent mode	2924
How to check the number of sessions that UTM proxy uses	2925
How to examine the firewall session list	2929
How to check wireless information	2930
How to verify connectivity to FortiGuard	2931
How to perform a sniffer trace (CLI and packet capture)	2931
How to debug the packet flow	2934
Troubleshooting resources	2936
Technical documentation	2936
Fortinet video library	2936
Release notes	2936
Knowledge base	2936
Fortinet technical discussion forums	2936
Fortinet training services online campus	2936
Fortinet customer support	2936
Chapter 28 - Virtual Domains	2937
What's new in FortiOS 6.0	2937

VDOMs overview.....	2938
Benefits of virtual domains.....	2938
Enabling and accessing VDOMs.....	2939
Configuring additional VDOMs.....	2943
VDOMs in NAT mode.....	2949
Using a VDOM in NAT/route mode.....	2949
Example configuration: VDOM in NAT/route mode.....	2953
VDOMs in transparent mode.....	2964
Transparent mode overview.....	2964
Operation mode differences in VDOMs.....	2965
Using a VDOM in transparent mode.....	2966
Example configuration: VDOM in transparent mode.....	2967
Security profiles and VDOMs.....	2981
VDOM-level security profiles.....	2981
Global security profiles.....	2981
Inter-VDOM routing.....	2982
Benefits of inter-VDOM routing.....	2982
Inter-VDOM configurations.....	2983
Configuring VDOM links.....	2986
Dynamic routing over inter-VDOM links.....	2989
HA virtual clusters and VDOM links.....	2990
Example configuration: inter-VDOM routing.....	2991
Troubleshooting VDOMs.....	3014
VDOM admin having problems gaining access.....	3014
FortiGate running very slowly.....	3014
General VDOM tips and troubleshooting.....	3015
Virtual FortiOS (Private Cloud Administration Guide).....	3019
What's new in virtual FortiOS 6.0.1.....	3019
What's new in virtual FortiOS 6.0.....	3019
Virtual FortiOS overview.....	3021
FortiGate VM models and licensing.....	3021
Registering FortiGate VM.....	3022
Downloading the FortiGate VM deployment package.....	3022
Deployment package contents.....	3023
Deploying the FortiGate VM appliance.....	3025
Performance and optimization.....	3025
Other virtual FortiOS products.....	3026
Deployment example – VMware.....	3037
Open the FortiGate VM OVF file with the vSphere client.....	3037
Configure FortiGate VM hardware settings.....	3042
Transparent mode VMware configuration.....	3043
High availability VMware configuration.....	3043

Power on your FortiGate VM.....	3044
Deployment example – MS Hyper-V.....	3045
Create the FortiGate VM virtual machine.....	3045
Configure FortiGate VM hardware settings.....	3052
High availability Hyper-V configuration.....	3063
Start the FortiGate VM.....	3064
Deployment example – KVM.....	3065
Create the FortiGate VM virtual machine.....	3065
Configure FortiGate VM hardware settings.....	3067
Start the FortiGate VM.....	3068
Deployment example – Open Xen.....	3069
Create the FortiGate VM virtual machine.....	3069
Deployment example – Citrix XenServer.....	3075
Create the FortiGate VM virtual machine (XenCenter).....	3075
Configure virtual hardware.....	3078
FortiGate VM initial configuration.....	3082
Set FortiGate VM port1 IP address.....	3082
Connect to the FortiGate VM web-based manager.....	3085
Upload the FortiGate VM license file.....	3085
Validate the FortiGate VM license with FortiManager.....	3087
Configure your FortiGate VM.....	3089
Chapter 30 - VoIP Solutions: SIP.....	3090
What's new in FortiOS 6.0.1.....	3091
What's new in FortiOS 6.0.....	3091
Inside FortiOS: Voice over IP (VoIP) protection.....	3092
Advanced voice over IP protection.....	3092
Carrier grade protection.....	3094
NAT/NAPT.....	3094
SIP ALG activation.....	3096
SIP over IPv6.....	3097
Platform support and hardware acceleration.....	3097
Common SIP VoIP configurations.....	3098
Peer to peer configuration.....	3098
SIP proxy server configuration.....	3098
SIP redirect server configuration.....	3099
SIP registrar configuration.....	3100
SIP with a FortiGate.....	3101
SIP messages and media protocols.....	3104
Hardware accelerated RTP processing.....	3106
SIP request messages.....	3106
SIP response messages.....	3107
SIP message start line.....	3109

SIP headers	3109
The SIP message body and SDP session profiles	3111
Example SIP messages	3114
The SIP session helper	3116
SIP session helper configuration overview	3116
Viewing, removing, and adding the SIP session helper configuration	3117
Changing the port numbers that the SIP session helper listens on	3117
Configuration example: SIP session helper in transparent mode	3118
SIP session helper diagnose commands	3121
The SIP ALG	3122
Enabling VoIP support from the GUI	3124
SIP ALG configuration overview	3124
VoIP profiles	3124
Changing the port numbers that the SIP ALG listens on	3126
Disabling the SIP ALG in a VoIP profile	3126
SIP ALG diagnose commands	3126
Conflicts between the SIP ALG and the session helper	3128
Stateful SIP tracking, call termination, and session inactivity timeout	3129
Adding a media stream timeout for SIP calls	3129
Adding an idle dialog setting for SIP calls	3130
Changing how long to wait for call setup to complete	3130
SIP and RTP/RTCP	3131
How the SIP ALG creates RTP pinholes	3132
Configuration example: SIP in transparent mode	3134
General configuration steps	3134
Configuration steps - GUI	3134
Configuration steps - CLI	3136
RTP enable/disable (RTP bypass)	3137
Opening and closing SIP register, contact, via and record-route pinholes	3138
Accepting SIP register responses	3139
How the SIP ALG performs NAT	3140
SIP ALG source address translation	3140
SIP ALG destination address translation	3141
SIP call re-invite messages	3141
How the SIP ALG translates IP addresses in SIP headers	3141
How the SIP ALG translates IP addresses in the SIP body	3144
SIP NAT scenario: source address translation (source NAT)	3145
SIP NAT scenario: destination address translation (destination NAT)	3147
SIP NAT configuration example: source address translation (source NAT)	3150
SIP NAT configuration example: destination address translation (destination NAT)	3152
SIP and RTP source NAT	3156
SIP and RTP destination NAT	3156

Source NAT with an IP pool.....	3158
Different source and destination NAT for SIP and RTP.....	3158
NAT with IP address conservation.....	3159
Controlling how the SIP ALG NATs SIP contact header line addresses.....	3160
Controlling NAT for addresses in SDP lines.....	3161
Translating SIP session destination ports.....	3161
Translating SIP sessions to multiple destination ports.....	3163
Adding the original IP address and port to the SIP message header after NAT.....	3164
Enhancing SIP pinhole security.....	3165
Hosted NAT traversal.....	3168
Configuration example: Hosted NAT traversal for calls between SIP Phone A and SIP Phone B.....	3169
Hosted NAT traversal for calls between SIP Phone A and SIP Phone C.....	3173
Restricting the RTP source IP.....	3173
SIP over IPv6.....	3174
Deep SIP message inspection.....	3175
Actions taken when a malformed message line is found.....	3176
Logging and statistics.....	3176
Deep SIP message inspection best practices.....	3176
Configuring deep SIP message inspection.....	3177
Blocking SIP request messages.....	3180
SIP rate limiting.....	3182
Limiting the number of SIP dialogs accepted by a security policy.....	3184
SIP logging.....	3185
Inspecting SIP over SSL/TLS (secure SIP).....	3186
Adding the SIP server and client certificates.....	3188
Adding SIP over SSL/TLS support to a VoIP profile.....	3188
SIP and HA—session failover and geographic redundancy.....	3190
SIP geographic redundancy.....	3190
Supporting geographic redundancy when blocking OPTIONS messages.....	3191
Support for RFC 2543-compliant branch parameters.....	3192
SIP and IPS.....	3193
SIP debugging.....	3194
SIP debug log format.....	3194
SIP-proxy filter per VDOM.....	3195
SIP-proxy filter command.....	3195
SIP debug setting.....	3195
Display SIP rate-limit data.....	3196
Supported RFCs.....	3197
What's new in FortiOS 6.0.1.....	3197
What's new in FortiOS 6.0.....	3197
Supported RFCs.....	3198

BGP.....	3198
Cryptography.....	3198
DHCP.....	3199
Diffserv.....	3199
DNS.....	3199
ICMP.....	3200
IP.....	3200
IP multicast.....	3200
IPsec.....	3201
IPv4.....	3201
IPv6.....	3201
IS-IS.....	3202
LDAP.....	3202
MPLS.....	3202
NAT.....	3202
OSPF.....	3203
PPP.....	3203
RADIUS.....	3203
RIP.....	3203
SIP.....	3203
SNMP.....	3204
SSL.....	3204
TCP.....	3204
TLS.....	3204
VPN.....	3205
Other protocols.....	3205
Miscellaneous.....	3205

Chapter 1 - What's New for FortiOS 6.0

Fortinet Security Fabric

This section introduces new Security Fabric features in FortiOS 6.0.

Security Fabric automation

User-defined Automations allow you to improve response times to security events by automating the activities between devices in the Security Fabric. You can monitor events from any source in the Security Fabric and set up action responses to any destination. To create an Automation, you can set up a Trigger event and response Actions that cause the Security Fabric to respond in a predetermined way. From the root FortiGate, you can set up event triggers for the following event types: compromised host, event log, reboot, conserve mode, high CPU, license expiry, High Availability (HA) failover, and configuration changes. The workflows have the means to launch the following actions in response: email, FortiExplorer notification, AWS Lambda and webhook. Additional actions are available for compromised hosts, such as: access layer quarantine, quarantine FortiClient via EMS, and IP ban.

FortiOS 6.0.2 adds the ability to test automation stitches using the `diagnose automation test` command.

For more information, see [Using the Fortinet Security Fabric on page 2337](#).

Security rating

The Security Rating feature (previously called the Security Fabric Audit) includes new security checks that can help you make improvements to your organization's network, such as enforce password security, apply recommended login attempt thresholds, encourage two factor authentication, and more.

For more information, see the [Fortinet Recommended Security Best Practices](#) document.

Security rating FortiGuard service

Security Rating is now a subscription service that FortiGuard offers when you purchase a Security Rating license. This service allows you to:

- Dynamically receive updates from FortiGuard.
- Run Security Rating checks for each licensed device in a Security Fabric.
- Run Security Rating checks in the background or on demand.
- Submit rating scores to FortiGuard and receive rating scores from FortiGuard, for ranking customers by percentile.

For more information, see [Using the Fortinet Security Fabric on page 2337](#).

Solution and service integration

In FortiOS 6.0, the Security Fabric extends to include more Fortinet products.

Wireless user quarantine

When you create or edit an SSID, you can enable the **Quarantine Host** option to quarantine devices that are connected in Tunnel-mode. The option to quarantine a device is available from the **Topology** and **FortiView** WiFi pages.

When a host is put into the quarantine VLAN, it gets an IP address from the quarantine VLAN DHCP server, and becomes part of the quarantined network.

For more information, see [Features for high-density deployments on page 1131](#).

Fortinet products can join the Security Fabric by serial number

Fortinet products can now easily and securely join the Security Fabric using an authorized device serial number.

For more information, see [Using the Fortinet Security Fabric on page 2337](#).

FortiMail integration

You can now add a FortiMail stats widget to the FortiGate Dashboard page to show mail detection stats from FortiMail. Other FortiMail integrations include the following:

- A FortiMail section that displays the FortiMail name, IP address, login and password is now available in the Security Fabric Settings page.
- FortiMail is now shown as a node in the topology tree view in the Fabric Settings page and in the Physical Topology and Logical Topology views.
- The topology views now show the number of FortiMail devices in the Security Fabric in the device summary.

For more information, see [Using the Fortinet Security Fabric on page 2337](#).

Synchronize the FortiManager IP address among all Security Fabric members

When you add a FortiManager to the root FortiGate of the Security Fabric, its configuration is now automatically synchronized with all devices in the Security Fabric. Central management features are now configured from the Security Fabric Settings page.

For more information, see [Using the Fortinet Security Fabric on page 2337](#).

Improve FortiAP and FortiSwitch support in Security Fabric views

The Security Fabric widget on the dashboard and the Security Fabric Settings page now show the FortiAP and FortiSwitch devices in the Security Fabric.

- You can now use new shortcuts to easily authorize any newly discovered devices and manage them.
- Switch stacking is now supported in the Physical and Logical topology views, and Inter-switch Link (ISL-LAG) is now identified by a thicker single line.

For more information, see [Using the Fortinet Security Fabric on page 2337](#).

EMS server support in Security Fabric topology

The FortiClient Endpoint Management System (EMS) can be enabled in FortiClient Endpoint profiles. This feature allows you to maintain FortiClient endpoint protection from FortiClient EMS and dynamically push configuration changes from the EMS to FortiClient endpoints. EMS server support is also integrated with Security Fabric Automation.

For more information, see [Using the Fortinet Security Fabric on page 2337](#).

Multi-cloud support (Security Fabric connectors)

Security Fabric multi-cloud support adds Security Fabric connectors to the Security Fabric configuration. Security Fabric connectors allow you to integrate Application Centric Infrastructure (ACI), Amazon Web Services (AWS), Microsoft Azure, VMware NSX, and Nuage Virtualized Services Platform configurations into the Security Fabric.

Additionally Cloud init support for Azure is now native to the cloud. FortiGate VM for Azure also supports bootstrapping.

For more information, see [Fabric Connectors on page 2362](#) and [Other virtual FortiOS products on page 3026](#).

Azure regional support

The Azure Security Fabric connector supports connecting to regional Azure public clouds. This change allows organizations in different regions to connect to their regional Azure public cloud if required for compliance or performance reasons.

For more information, see [Fabric Connectors on page 2362](#) and [Other virtual FortiOS products on page 3026](#).

GUI change for single sign-on configuration

In FortiOS 6.0.1, the options to configure single sign-on in the FortiGate GUI are now located in the **Security Fabric > Fabric Connectors** menu.

Manageability

This section introduces new manageability features in FortiOS 6.0.

Asset tagging

You can use the new Asset Tagging system to create tags to separate and categorize network objects, interfaces, and devices. Tags are flexible, easy to configure, and useful for comprehensive monitoring, audit reporting, and more.

For more information, see [Multi-dimension tagging on page 2779](#).

FortiSwitch network assisted device detection and destination name resolution

Device detection now extends to managed FortiSwitches since some devices may not be visible to the FortiGate that manages them. Devices that are connected to a FortiSwitch are more visible to the FortiGate that manages them and to the Security Fabric.

FortiSwitch destination name resolution clearly presents destination objects and the aggregation of related IP addresses with domains. It also applies Internet Service Database (ISDB) mapping for destination data.

For more information, see [Managing “bring your own device” on page 1892](#) and [FortiLink configuration using the FortiGate GUI on page 1921](#).

Global security profiles

Global Security Profiles can be used by multiple VDOMs instead of creating identical profiles for each VDOM. You can create global security profiles for the following security features:

- Antivirus
- Application control
- Data leak prevention
- Intrusion protection
- Web filtering

For more information, see [Global security profiles on page 2981](#).

Networking

This section introduces new networking features in FortiOS 6.0.

SD-WAN improvements

FortiOS 6.0 introduces the following SD-WAN features:

- Multiple server support for health checks
- Internet service groups
- Bandwidth options in SD-WAN rules
- Custom profiles in SD-WAN rules
- DSCP tagging of forwarded packets in SD-WAN rules

For more information, see [SD-WAN on page 2273](#).

Multipath intelligence and performance SLAs

SD-WAN performance Service-Level Agreements (SLAs) incorporate multilayer SLA monitoring of link selection. To help handle emergency load or outages you can select links based on weight and SLA priority and then return to defaults once the network stabilizes. Also, traffic shaping and application intelligence have been added to the SD-WAN configuration, which gives you more control of SD-WAN traffic.

For more information, see [SD-WAN on page 2273](#).

Application awareness

You can now use application control and application control group options in SD-WAN rules.

Internet Service support is also increased from a single Internet Service to Internet Service groups.

For more information, see [Application control settings in SD-WAN rules on page 2273](#).

BGP dynamic routing and IPv6 support for SD-WAN

FortiOS 6.0 introduces support for dynamic router for an SD-WAN configuration. You can set up a route map and add a route tag to the route map. Then, you can create an SD-WAN configuration, a health check, and a service for it. When you create the service, you add the configured route tag that you created in the route map to the service.

For more information, see [SD-WAN on page 2273](#).

Interface-based traffic shaping

In FortiOS 6.0, you can now enable traffic shaping on an interface. Interface-based traffic shaping allows you to enforce bandwidth limits by traffic type for individual interfaces.

For more information, see [Interface-based traffic shaping on page 2805](#).

Cloud-assisted one-click VPN

One-click VPN (OCVPN) is a cloud-based solution that greatly simplifies the provisioning and configuration of IPsec VPN. The administrator enables OCVPN with a single click, adds the required subnets, and then the configuration is complete. The OCVPN updates each FortiGate automatically as devices join and leave the VPN, as subnets are added and removed, when dynamic external IP addresses change (for example, DHCP or PPPoE), and when WAN interface bindings change (as in the case of dual WAN redundancy).

For more information, see [One-Click VPN \(OCVPN\) on page 1719](#).

IPv6 enhancements

The following new IPv6 features have been added.

- IPv6 captive portal
- IPv6 FQDN and wildcard firewall addresses
- IPv6 ISIS dynamic routing
- DHCPv6 server prefix delegation
- IPv6 DFD and VRRP

For more information, see [IPv6 on page 613](#).

NAT enhancements

The following new NAT features have been added.

- Central source NAT (SNAT) policies now include a comment field
- Port block allocation timeout is configurable
- NAT46 IP pools
- VRRP HA supports firewall virtual IPs (VIPs) and IP pools

For more information, see [NAT on page 559](#).

EMAC-VLAN support

The media access control (MAC) virtual local area network (VLAN) feature in Linux allows you to configure multiple virtual interfaces with different MAC addresses (and therefore different IP addresses) on a physical interface.

For more information, see [Enhanced MAC VLANs on page 2049](#).

Security

This section introduces new security features in FortiOS 6.0.

FortiGuard virus outbreak prevention

FortiGuard virus outbreak prevention is an additional layer of protection that keeps your network safe from newly emerging malware. Quick virus outbreaks can infect a network before signatures can be developed to stop them. Outbreak protection stops these virus outbreaks until signatures become available in FortiGuard.

For more information, see [FortiGuard virus outbreak prevention on page 2432](#).

FortiGuard content disarm and reconstruction

Content Disarm and Reconstruction (CDR) removes exploitable content and replaces it with content that's known to be safe. As files are processed through an enabled AntiVirus profile, content that's found to be malicious or unsafe is replaced with content that allows the traffic to continue, but doesn't put the recipient at risk.

Content that can be scanned includes PDF and Microsoft Office files leaving the network on CDR-supported protocols (such as, HTTP web download, SMTP email send, IMAP and POP3 email retrieval—MAPI isn't supported).

This feature works even if FortiSandbox is not configured, but only if you want to discard the original file. If FortiSandbox is configured and it responds that the file is clean, it passes the content unmodified.

For more information, see [Content Disarm and Reconstruction \(CDR\) on page 2429](#).

Application groups for NGFW policies

When a FortiGate operates in NGFW policy mode, you can create application groups when you add NGFW policies. Then, when you add IPv4 or IPv6 policies you can create application groups to simplify policy creation.

For more information, see [Application groups for NGFW policies on page 532](#).

Application control rule sequencing

To have more control over application control outcomes, you can control the order that application signatures appear in application control sensors. Signatures for applications that are more sensitive can appear higher in the list so they get matched first.

For more information, see [Application control on page 2485](#).

Threat Feeds (external dynamic block lists)

This feature introduces the ability to dynamically import external block lists from an HTTP server. You can use the block lists to enforce special security requirements that your organization has. This can include long term policies to always block access to some websites or short time requirements to block access to known compromised locations. Since the lists are dynamically imported any changes made to the list are instantly imported by FortiOS. Dynamic block lists can be added to:

- Web Filter profiles and SSL inspection exemptions.
- DNS Filter profiles and "Source/Destination" addresses in proxy policies.

In each profile, the administrator can configure multiple external block lists.

For more information, see [Threat Feed Connectors on page 2462](#).

FortiAP-S bridge mode security profiles

If you have enabled bridge mode for a managed FortiAP-S, you can add a UTM profile to the wireless controller configuration that allows you to apply the following security profile features to all traffic accepted by the managed FortiAP-S:

- AntiVirus (including Botnet protection),
- IPS,
- Application control, and
- Web Filtering.

For more information, see [FortiAP-S bridge mode security profiles on page 1144](#).

Chapter 2 - Getting Started

- [Installation](#) discusses installing a FortiGate in your network.
- [Using the GUI](#) highlights features of the graphical user interface (GUI).
- [Using the CLI](#) provides a high level overview of the command line interface (CLI) for FortiOS.
- [FortiExplorer for iOS](#) provides instructions for connecting to a FortiGate using the FortiExplorer for iOS app.
- [LED Specifications](#) presents a short guide to LED status indicators.
- [Inspection Mode](#) summarizes proxy-based and flow-based inspection modes.
- [Basic administration](#) explains basic tasks for setting up a new FortiGate and for updating firmware.
- [Troubleshooting your FortiGate installation](#) provides troubleshooting tips if your FortiGate installation is unsuccessful.
- [Resources](#) provides a list of documents to help you with more advanced FortiGate configurations.

Differences between models

Before you get started, note that not all FortiGate models have the same features. This is especially true of the desktop or entry-level models: FortiGate / FortiWiFi models 30 to 90. If you are using one of these FortiGate models, you may have some difficulties accessing certain features.

The entry-level, or desktop, models can connect to the internet in two simple steps. They also have a number of features that are only available using the CLI, rather than appearing in the GUI.

- [Quick installation using DHCP](#)
- [CLI-only features](#)



Consult your model's Quick Start Guide, [hardware manual](#), or the [Feature / Platform Matrix](#) for further information about features that vary by model.

The Fortinet Cookbook site has a section on [hardware](#) that provides how-to recipes and articles on features that are unique to certain models.

FortiGate models differ principally by the names used and the features available:

- Naming conventions may vary between FortiGate models. For example, on some models the hardware switch interface used for the local area network is called **lan**, while on other units it is called **internal**.
- Certain features are not available on all models. Additionally, a particular feature may be available only through the CLI on some models, while that same feature may be viewed in the GUI on other models.

If you believe your FortiGate model supports a feature that does not appear in the GUI, go to **System > Feature Visibility** and confirm that the feature is enabled. For more information, see [Feature Visibility on page 69](#).

What's new in FortiOS 6.0

The following list contains new getting started features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- [New dashboard widget for "Botnet Activity" on page 68](#)
- [Dashboard widget improvements for "Administrators" on page 67](#)
- ["Modifying dashboard widget titles" on page 68](#)

Installation

This section discusses how to install your FortiGate and use it in your network, after completion of the initial setup outlined in the FortiGate model's Quick Start Guide.

The following topics are included in this section:

- [Quick installation using DHCP](#)
- [Installing a FortiGate in NAT/Route mode](#)
- [Using a virtual wire pair](#)

Quick installation using DHCP

Most of the FortiGate desktop models have a default configuration that includes a DHCP server on the **lan** (or **internal**) interface and a security policy that securely allows all sessions from the internal network to reach the Internet. Because of this, you can connect your desktop FortiGate to the Internet in two simple steps:



Note that, in order to use this installation method, your ISP must provide connectivity with DHCP and accept DHCP requests without authentication. You must also use IPv4 to connect your FortiGate to the Internet.

1. Connect the **wan** interface on your FortiGate to your ISP-supplied equipment, and connect the internal network to the default **lan** interface on your FortiGate. Turn on the ISP's equipment, the FortiGate, and the computers on the internal network.
2. For computers on the internal network:
 - a. **Windows Vista/7/8/10 users:**
 - i. Go to **Network and Sharing Center** and select **Change adapter settings**.
 - ii. Open the **Local Area Connection** (Ethernet or WiFi, whichever applies) and select **Properties**.
 - iii. Select **Internet Protocol Version 4 (TCP/IPv4)** and then select **Properties**.
 - iv. Select **Obtain an IP address automatically** and **Obtain DNS server address automatically**.
 - v. Click **OK**.
 - b. **Mac OS X users:**
 - i. Go to **System Preferences > Network** and select your Ethernet connection.
 - ii. Set **Configure IPv4** to **Using DHCP**.

Results

To confirm successful Internet connectivity from any computer on the internal network, open a web browser and browse to any website.

Installing a FortiGate in NAT/route mode

There are two main ways to install a FortiGate using network address translation (NAT)/route mode: [Standard installation in NAT/route mode](#), where Internet access is provided by a single Internet service provider (ISP), and [Redundant Internet installation](#), where two ISPs are used.

NAT/Route mode vs. transparent mode

A FortiGate can operate in one of two modes: NAT/route or transparent.

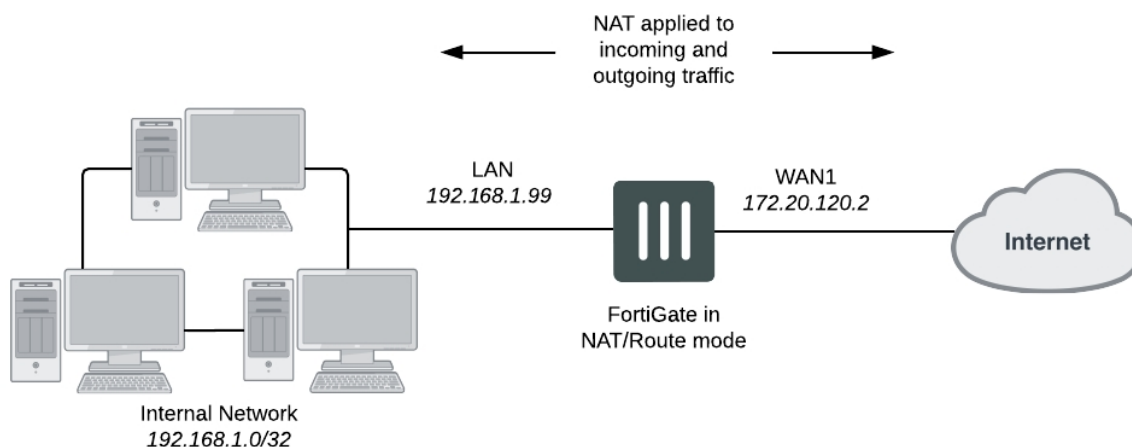
The most common of the two operating modes is NAT/route mode, where a FortiGate is installed as a gateway or router between two networks. In most cases, it is used between a private network and the Internet. This allows the FortiGate to hide the IP addresses of the private network using NAT. NAT/route mode is also used when two or more ISPs provide the FortiGate with redundant Internet connections.

A FortiGate in transparent mode is installed between the internal network and the router. In this mode, the FortiGate does not make any changes to IP addresses and only applies security scanning to traffic. When a FortiGate is added to a network in transparent mode, no network changes are required, except to provide the FortiGate with a management IP address. Transparent mode is used primarily when there is a need to increase network protection but changing the configuration of the network itself is impractical.

For more information about transparent mode, see the [Transparent Mode](#) chapter of the handbook.

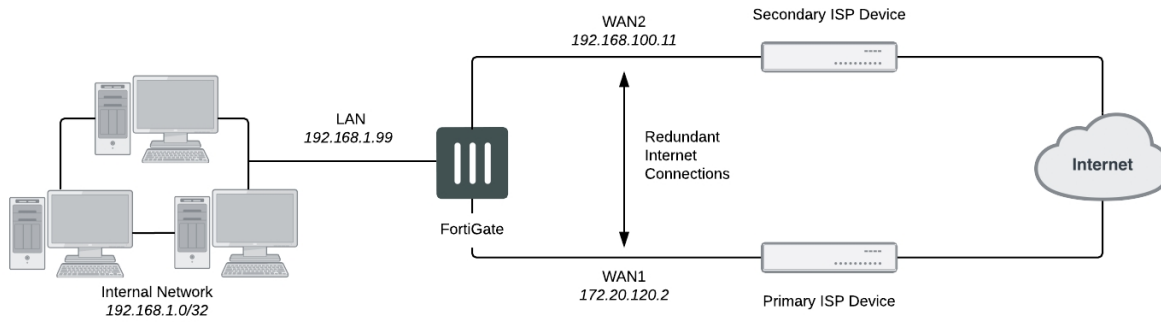
Standard installation in NAT/route mode

In this configuration, a FortiGate is installed as a gateway or router between a private network and the Internet. By using NAT mode, the FortiGate is able to hide the IP addresses of the private network.



Redundant Internet installation

In this configuration, a WAN link interface is created that provides the FortiGate with redundant Internet connections from two ISPs. The WAN link interface combines these two connections, allowing the FortiGate to treat them as a single interface.



Installing a FortiGate with redundant Internet



If you have previously configured your FortiGate using the standard installation, you will have to delete all routes and policies referring to an interface that will be used to provide redundant Internet. This includes the default Internet access policy that is included on many FortiGate models.

1. Connect your ISP devices to your FortiGate's Internet-facing interfaces (typically WAN1 and WAN2).
2. Go to **Network > Interfaces** to create a WAN link interface, which is used to group multiple Internet connections together so that the FortiGate can treat them as a single interface.
3. Set the interface **Status** to **Enable**.
4. Under **SD-WAN Interface Members**, click on the plus sign and then on the down arrow to open the dropdown menu. Select WAN1 as the **Interface** and enter the **Gateway** IP provided by your primary ISP. Do the same for WAN2, but use the Gateway IP provided by your secondary ISP.
5. Select an appropriate method for the **SD-WAN Usage** from the following options, and **Apply** your changes when finished:
 - **Bandwidth** - A bandwidth cap is defined for active members of the SD WAN link.
 - **Volume** - A volume ratio is set for each active member.
 - **Sessions** - A sessions ratio is set for each active member.
6. Go to **Network > Static Routes** and create a new default route. Set **Interface** to the SD-WAN link.
7. Go to **Policy & Objects > IPv4 Policy** and select **Create New** to add a security policy that allows users on the private network to access the Internet.


Using a virtual wire pair

A virtual wire pair consists of two interfaces that do not have IP addressing and are treated similar to a transparent mode VDOM. All traffic received by one interface in the virtual wire pair can only be forwarded out the other interface, provided that a virtual wire pair firewall policy allows this traffic. Traffic from other interfaces cannot be routed to the interfaces in a virtual wire pair.

Virtual wire pairs are useful for atypical topologies where MAC addresses do not behave normally. For example, port pairing can be used in a Direct Server Return (DSR) topology where the response MAC address pair may not match the request's MAC address pair.

Virtual wire pairing replaces the port pairing feature available in earlier firmware versions. Unlike port pairing, virtual wire pairing can be used for FortiGate in both NAT/Route and Transparent modes.

In the example configuration below, a virtual wire pair (consisting of port3 and port4) makes it easier to protect a web server that is behind a FortiGate operating as an Internal Segmentation Firewall (ISFW). Users on the internal network will access the web server through the ISFW over the virtual wire pair.



Adding a virtual wire pair and virtual wire pair policy



Interfaces used in a virtual wire pair cannot be used to access the ISFW FortiGate. Before creating a virtual wire pair, make sure you have a different port configured to allow admin access using your preferred protocol.

1. Go to **Network > Interfaces** and select **Create New > Virtual Wire Pair**.
2. Select the interfaces to add to the virtual wire pair. These interfaces cannot be part of a switch, such as the default **lan/internal** interface.
3. (Optional) If desired, enable **Wildcard VLAN**.
4. Select **OK**.
5. Go to **Policy & Objects > IPv4 Virtual Wire Pair Policy**, select the virtual wire pair, and select **Create New**.
6. Select the direction that traffic is allowed to flow.
7. Configure the other firewall options as desired.
8. Select **OK**.
9. If necessary, create a second virtual wire pair policy to allow traffic to flow in the opposite direction.



If you have a USB-wan interface, it will not be included in the interface list when building a wired-pair.

Results

Traffic can now flow through the FortiGate using the virtual wire pair. For more information on this feature, see the [Networking](#) chapter.

Using the GUI

This section presents an introduction to the graphical user interface (GUI) on your FortiGate, also called the web-based manager.

The following topics are included in this section:

- [Connecting to the GUI](#)
- [Menus](#)
- [Dashboard](#)
- [Feature Visibility](#)
- [Tables](#)
- [Text strings](#)

Connecting to the GUI using a web browser



The graphical user interface is best displayed using a 1280 x 1024 resolution. Check the [FortiOS Release Notes](#) for information about browser compatibility.

In order to connect to the GUI using a web browser, an interface must be configured to allow administrative access over HTTPS or over both HTTPS and HTTP. By default, an interface has already been set up that allows HTTPS access, with the IP address 192.168.1.99.

Browse to <https://192.168.1.99> and enter your username and password. If you have not changed the admin account's password, use the default user name, `admin`, and leave the password field blank.

The GUI will now be displayed in your browser.

If you wish to use a different interface to access the GUI, do the following:

1. Go to **Network > Interfaces** and edit the interface you wish to use for access. Take note of its assigned IP address.
2. Beside **Administrative Access**, select **HTTPS**, and any other protocol you require. You can also select **HTTP**, although this is not recommended as the connection will be less secure.
3. Select **OK**.
4. Browse to the IP address using your chosen protocol.

Results

The GUI will now be displayed in your browser.

Menus



If you believe your FortiGate model supports a menu that does not appear in the GUI as expected, go to **System > Feature Visibility** and ensure the feature is enabled. For more information, see ["Feature Visibility" on page 1](#).

The GUI contains the following main menus, which provide access to configuration options for most FortiOS features:

Dashboard	<p>The dashboard displays various widgets that display important system information and allow you to configure some system options.</p> <p>For more information, see "Dashboard" on page 1.</p>
Security Fabric	<p>Access the physical topology, logical topology, audit, and settings features of the Fortinet Security Fabric.</p> <p>For more information, see the Fortinet Security Fabric handbook.</p>
FortiView	<p>A collection of dashboards and logs that give insight into network traffic, showing which users are creating the most traffic, what sort of traffic it is, when the traffic occurs, and what kind of threat the traffic may pose to the network.</p> <p>For more information, see the FortiView handbook.</p>
Network	<p>Options for networking, including configuring system interfaces and routing options.</p> <p>For more information, see the Networking handbook.</p>
[[[Undefined variable FortiOSGUIVariables.Kee]]]	<p>Configure system settings, such as administrators, FortiGuard, and certificates.</p> <p>For more information, see the System Administration handbook.</p>
Policy & Objects	<p>Configure firewall policies, protocol options, and supporting content for policies, including schedules, firewall addresses, and traffic shapers.</p> <p>For more information, see the Firewall handbook.</p>
Security Profiles	<p>Configure your FortiGate's security features, including AntiVirus, Web Filtering, and Application Control.</p> <p>For more information, see the Security Profiles handbook.</p>
VPN	<p>Configure options for IPsec and SSL virtual private networks (VPNs).</p> <p>For more information, see the IPsec VPN and SSL VPN handbooks.</p>

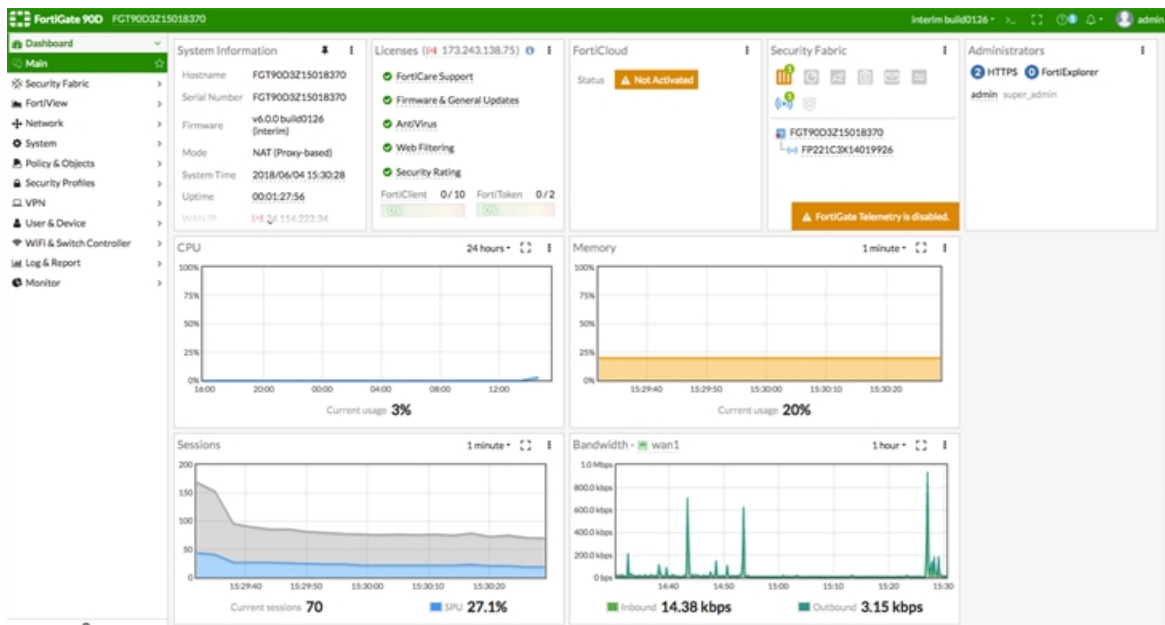
User & Device	Configure user accounts, groups, and authentication methods, including external authentication and single sign-on (SSO).
WiFi & Switch Controller	<p>Configure the unit to act as a wireless network controller, managing the wireless Access Point (AP) functionality of FortiWiFi and FortiAP units.</p> <p>On certain FortiGate models, this menu has additional features allowing for FortiSwitch units to be managed by the FortiGate.</p> <p>For more information, see the FortiWiFi and FortiAP Configuration Guide.</p>
Log & Report	<p>Configure logging and alert email as well as reports.</p> <p>For more information, see the Logging and Reporting handbook.</p>
Monitor	View a variety of monitors, including the Routing Monitor, VPN monitors for both IPsec and SSL, monitors relating to wireless networking, and more.

Dashboard

The FortiOS **Dashboard** consists of a Network Operations Center (NOC) view with a focus on alerts. Widgets are interactive. By clicking or hovering over most widgets, the user can see additional information or follow links to other pages.

The dashboard and its widgets include:

- Multiple dashboard support
- VDOM and global dashboards
- Widget resize control
- Notifications on the top header bar



The following widgets are displayed by default:

Widget	Description
System Information	The System Information widget lists information relevant to the FortiGate system, including hostname, serial number, and firmware.
Security Fabric	The Security Fabric widget displays a visual summary of many of the devices in the Security Fabric. For more information, see the Security Fabric handbook.
CPU	The real-time CPU usage is displayed for different timeframes.

Widget	Description
Licenses	<p>Hovering over the Licenses widget results in the display of status information (and, where applicable, database information) on the licenses for FortiCare Support, Firmware & General Updates, AntiVirus, Web Filtering, Security Rating, FortiClient, and FortiToken. Note that Mobile Malware is not a separate service in FortiOS 6.0.0. The Mobile Malware subscription is included with the AntiVirus subscription.</p> <p>Clicking in the Licenses widget provides you with links to other pages, such as System > FortiGuard or contract renewal pages.</p>
FortiCloud	This widget displays FortiCloud status and provides a link to activate FortiCloud.
Administrators	<p>This widget allows you to view:</p> <ul style="list-style-type: none"> • which administrators are logged in and how many sessions are active (a link directs you to a page displaying active administrator sessions) • all connected administrators and the protocols used by each
Memory	Real-time memory usage is displayed for different time frames. Hovering over any point on the graph displays percentage of memory used along with a timestamp.
Sessions	<p>Hovering over the Sessions widget allows you to view memory usage data over time. Click on the down arrow to change the timeframe displayed.</p> <p>Security processing unit, or SPU, percentage is displayed if your FortiGate includes an SPU. Likewise, nTurbo percentage is displayed if supported by your FortiGate. See the Hardware Acceleration chapter for details.</p>
Bandwidth	<p>Hover over the Bandwidth widget to display bandwidth usage data over time. Click on the down arrow to change the timeframe displayed. Bandwidth is displayed for both incoming and outgoing traffic.</p>
Virtual Machine	<p>The VM widget (shown by default in the dashboard of a FortiOS VM device) includes:</p> <ul style="list-style-type: none"> • License status and type • CPU allocation usage • License RAM usage • VMX license information (if the VM supports VMX) <p>If the VM license specifies 'unlimited' the progress bar is blank. If the VM is in evaluation mode, it is yellow (warning style) and the dashboard shows the number of evaluation days used.</p>

The following optional widgets are also available:

- FortiView
- Host Scan Summary

- Vulnerabilities Summary
- Botnet Activity
- HA Status
- Log Rate
- Session Rate
- Security Fabric Score
- Advanced Threat Protection Statistics
- Interface Bandwidth

Modifying dashboard widget titles

Dashboard widget titles can be modified so that widgets with different filters applied can be easily differentiated. The widget has a default title unless you set a new title.

Syntax

```
config system admin
  edit <name>
    config gui-dashboard
      config widget
        edit 9
          set type fortiview
          ...
          set title "test source by bytes"
        end
      end
    end
  end
```

Feature Visibility

Feature Visibility is used to control which features are visible in the GUI. This allows you to hide features that are not being used. Some features are also disabled by default and must be enabled in order to configure them through the GUI.

Feature Visibility only alters the visibility of these features, rather than their functionality. For example, disabling web filtering on the **Feature Visibility** page does not remove web filtering from the FortiGate, but removes the option of configuring web filtering from the GUI. Configuration options will still be available using the CLI.

Enabling/disabling features

Feature Visibility can be found at **System > Feature Visibility**. Ensure that all features you wish to configure in the GUI are turned on, and that features you wish to hide are turned off. When you have finished, select **Apply**.

Security feature presets

The main security features can be toggled individually, however six system presets (or **Feature Sets**) are available:

- **NGFW** should be chosen for networks that require application control and protection from external attacks.
- **ATP** should be chosen for networks that require protection from viruses and other external threats.
- **WF** should be chosen for networks that require web filtering.
- **NGFW + ATP** should be chosen for networks that require protection from external threats and attacks.
- **UTM** should be chosen for networks that require protection from external threats and wish to use security features that control network usage. This is the default setting.
- **Custom** should be chosen for networks that require customization of available features (including the ability to select all features).

Tables

Many of the GUI pages contain tables of information that you can filter to display specific information. Administrators with read and write access can define the filters.

Navigation

Some tables contain information and lists that span multiple pages. Navigation controls appear at the bottom of the page.

Filters

Filters are used to locate a specific set of information or content within multiple pages. These are especially useful in locating specific log entries. The specific filtering options vary, depending on the type of information in the log.

To create a filter, select **Add Filter** at the top of the page. A list of the available fields for filtering will be shown.

Column settings

Column settings are used to select the types of information displayed on a certain page. Some pages have large amounts of information available and not all content can be displayed on a single screen. Some pages may even contain content that is irrelevant to you. Using column settings, you can choose to display only relevant content.

To view configure column settings, right-click the header of a column and select the columns you wish to view and deselect any you wish to hide. After you have finished making your selections, click **Apply** (you may need to scroll down the list to do so).

Any changes that you make to the column settings are stored in the unit's configuration. To return columns to the default state for any given page, right-click any header and select **Reset Table**.

Copying objects

In tables containing configuration objects, such as the policy table found at **Policy & Objects > IPv4 Policy**, you have the option to copy an object. This allows you to create a copy of that object, which you can then configure as needed. You can also reverse copy a policy to change the direction of the traffic impacted by that policy.

To copy an object:

1. Select that object, then right-click to make a menu appear and select the **Copy** option.
2. Right-click the row in the table that is either above or below where you want the copied object to be placed, select the **Paste** option and indicate **Above** or **Below**.

Reverse cloning works much the same way. Instead of selecting **Copy**, select **Clone Reverse**.

Once the policy is copied, you must give it a name, configure as needed, and enable it.

Editing objects

Some tables allow you to edit parts of the configuration directly on the table itself. For example, security features can be added to an existing firewall policy from the policy list by clicking on the plus sign in the **Security Profiles** column and selecting the desired profiles.

If this option is not immediately available, check to see that the column is not hidden (see [Column settings](#)). Otherwise, you must select the object and open the policy by selecting the **Edit** option found at the top of the page.

Text strings

The configuration of a FortiGate is stored in the FortiOS configuration database. To change the configuration, you can use the GUI or CLI to add, delete, or change configuration settings. These changes are stored in the database as you make them. Individual settings in the configuration database can be text strings, numeric values, selections from a list of allowed options, or on/off (enable/disable) settings.

Entering text strings (names)

Text strings are used to name entities in the configuration. For example, the name of a firewall address, the name of an administrative user, and so on. You can enter any character in a FortiGate configuration text string, except the following characters that present cross-site scripting (XSS) vulnerabilities:

- “ (double quote)
- & (ampersand)
- ' (single quote)
- < (less than)
- > (greater than)

Most GUI text string fields make it easy to add an acceptable number of characters and prevent you from adding the XSS vulnerability characters.



There is a different character limitation for VDOM names and hostnames. The only valid characters are numbers (0-9), letters (a-z, A-Z), and special characters - (dash) and _ (underscore).

You can also use the `tree` command in the CLI to view the number of characters allowed in a name field. For example, firewall address names can contain up to 64 characters. When you add a firewall address to the GUI, you are limited to entering 64 characters in the firewall address name field. From the CLI you can enter the following `tree` command to confirm that the firewall address `name` field allows 64 characters.

```
config firewall address
  tree
  -- [address] --*name (64)
  |- uuid
  |- subnet
  |- type
  |- start-ip
  |- end-ip
  |- fqdn (256)
  |- country (3)
  |- cache-ttl (0,86400)
  |- wildcard
  |- comment
  |- visibility
  |- associated-interface (36)
  |- color (0,32)
  |- [tags] --*name (65)
  +- allow-routing
```

The `tree` command output also shows the number of characters allowed for other firewall address name settings. For example, the fully qualified domain name (`fqdn`) field can contain up to 256 characters.

Entering numeric values

Numeric values set various sizes, rates, addresses, and other numeric values (e.g. a static routing priority of 10, a port number of 8080, an IP address of 10.10.10.1). Numeric values can be entered as a series of digits without spaces or commas (for example, 10 or 64400), in dotted decimal format (for example the IP address 10.10.10.1) or, as in the case of MAC or IPv6 addresses, separated by colons (e.g. the MAC address 00:09:0F:B7:37:00). Most numeric values are standard base 10 numbers, but some fields, such as MAC addresses, require hexadecimal numbers.

Most GUI numeric value fields make it easy to add the acceptable number of digits within the allowed range. CLI help text includes information about allowed numeric value ranges. Both the GUI and the CLI prevent you from entering invalid numbers.

Using the CLI

The command line interface (CLI) is an alternative configuration tool to the GUI or web-based manager. While the configuration of the GUI uses a point-and-click method, the CLI requires typing commands or uploading batches of commands from a text file, like a configuration script.

This section explains common CLI tasks that an administrator performs on a regular basis and includes the topics:

- [Connecting to the CLI](#)
- [Command syntax](#)
- [Sub-commands](#)
- [Permissions](#)
- [Tips](#)

Connecting to the CLI

You can access the CLI in three ways:

- [Locally with a console cable](#) — Connect your computer directly to the console port of your FortiGate. Local access is required in some cases:
 - If you are installing your FortiGate for the first time and it is not yet configured to connect to your network, you may only be able to connect to the CLI using a local serial console connection, unless you reconfigure your computer's network settings for a peer connection.
 - Restoring the firmware utilizes a boot interrupt. Network access to the CLI is not available until after the boot process has completed, making local CLI access the only viable option.
- [Through the network](#) — Connect your computer through any network interface attached to one of the network ports on your FortiGate. The network interface must have enabled Telnet or SSH administrative access if you connect using an SSH/Telnet client, or HTTP/HTTPS administrative access if you connect by accessing the **CLI Console** in the GUI. The CLI console can be accessed from the upper-right hand corner of the screen and appears as a slide-out window.
- [Locally with FortiExplorer for iOS](#) — Use the FortiExplorer app on your iOS device to configure, manage, and monitor your FortiGate.

Connecting to the CLI using a local console

Local console connections to the CLI are formed by directly connecting your management computer or console to the FortiGate unit, using its DB-9 or RJ-45 console port. To connect to the local console you need:

- A computer with an available serial communications (COM) port.
- The RJ-45-to-DB-9 or null modem cable included in your FortiGate package.
- Terminal emulation software such as HyperTerminal for Microsoft Windows.

The following procedure describes the connection using Microsoft HyperTerminal software; steps may vary with other terminal emulators.

To connect to the CLI using a local serial console connection

1. Using the null modem or RJ-45-to-DB-9 cable, connect the FortiGate unit's console port to the serial communications (COM) port on your management computer.
2. On your management computer, start HyperTerminal.
3. For the **Connection Description**, enter a **Name** for the connection, and select **OK**.
4. On the **Connect using** drop-down, select the communications (COM) port on your management computer you are using to connect to the FortiGate unit.
5. Select **OK**.
6. Select the following **Port** settings and select **OK**.

Bits per second	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	None

7. Press **Enter** or **Return** on your keyboard to connect to the CLI.
8. Type a valid administrator account name (such as `admin`) and press **Enter**.
9. Type the password for that administrator account and press **Enter**. (In its default state, there is no password for the `admin` account.)

The CLI displays the following text:

```
Welcome!
Type ? to list available commands.
```

You can now enter CLI commands, including configuring access to the CLI through SSH or Telnet.

Enabling access to the CLI through the network (SSH or Telnet)

SSH or Telnet access to the CLI is accomplished by connecting your computer to the FortiGate unit using one of its RJ-45 network ports. You can either connect directly, using a peer connection between the two, or through any intermediary network.



If you do not want to use an SSH/Telnet client and you have access to the web-based manager, you can alternatively access the CLI through the network using the **CLI Console** widget in the web-based manager.

You must enable SSH and/or Telnet on the network interface associated with that physical network port. If your computer is not connected directly or through a switch, you must also configure the FortiGate unit with a static route to a router that can forward packets from the FortiGate unit to your computer. You can do this using either a local console connection or the web-based manager.

Requirements

- A computer with an available serial communications (COM) port and RJ-45 port
- Terminal emulation software such as HyperTerminal for Microsoft Windows

- The RJ-45-to-DB-9 or null modem cable included in your FortiGate package
- A network cable
- Prior configuration of the operating mode, network interface, and static route.

To enable SSH or Telnet access to the CLI using a local console connection

1. Using the network cable, connect the FortiGate unit's network port either directly to your computer's network port, or to a network through which your computer can reach the FortiGate unit.
2. Note the number of the physical network port.
3. Using a local console connection, connect and log into the CLI.
4. Enter the following command:

```
config system interface
  edit <interface_str>
    set allowaccess <protocols_list>
  end
```

where:

- <interface_str> is the name of the network interface associated with the physical network port and containing its number, such as `port1`.
- <protocols_list> is the complete, space-delimited list of permitted administrative access protocols, such as `https ssh telnet`.

For example, to exclude HTTP, HTTPS, SNMP, and PING, and allow only SSH and Telnet administrative access on `port1`, enter the following:

```
config system interface
  edit port1
    set allowaccess ssh telnet
  end
```

5. To confirm the configuration, enter the command to display the network interface's settings.

```
show system interface <interface_str>
```

The CLI displays the settings, including the allowed administrative access protocols, for the network interfaces.

Connecting to the CLI using SSH

Once the FortiGate unit is configured to accept SSH connections, you can use an SSH client on your management computer to connect to the CLI.

Secure Shell (SSH) provides both secure authentication and secure communications to the CLI. FortiGate units support 3DES and Blowfish encryption algorithms for SSH.

Before you can connect to the CLI using SSH, you must first configure a network interface to accept SSH connections. The following procedure uses PuTTY. Steps may vary with other SSH clients.

To connect to the CLI using SSH

1. On your management computer, start an SSH client.
2. In **Host Name (or IP address)**, enter the IP address of a network interface on which you have enabled SSH administrative access.
3. Set a **Port** of 22.

4. For the **Connection type**, select **SSH**.

5. Select **Open**.

The SSH client connects to the FortiGate unit.

The SSH client may display a warning if this is the first time you are connecting to the FortiGate unit and its SSH key is not yet recognized by your SSH client, or if you have previously connected to the FortiGate unit but used a different IP address or SSH key. This is normal if your management computer is directly connected to the FortiGate unit with no network hosts between them.

6. Click **Yes** to verify the fingerprint and accept the FortiGate unit's SSH key. You will not be able to log in until you have accepted the key.
7. The CLI displays a login prompt.
8. Type a valid administrator account name (such as `admin`) and press **Enter**.
9. Type the password for this administrator account and press **Enter**.

The FortiGate unit displays a command prompt (its hostname followed by a #). You can now enter CLI commands.



If three incorrect login or password attempts occur in a row, you will be disconnected. If this occurs, wait one minute, then reconnect to attempt the login again.

Connecting to the CLI using Telnet

Once the FortiGate unit is configured to accept Telnet connections, you can use a Telnet client on your management computer to connect to the CLI.



Telnet is not a secure access method. SSH should be used to access the CLI from the Internet or any other untrusted network.

Before you can connect to the CLI using Telnet, you must first configure a network interface to accept Telnet connections.

To connect to the CLI using Telnet

1. On your management computer, start a Telnet client.
2. Connect to a FortiGate network interface on which you have enabled Telnet.
3. Type a valid administrator account name (such as `admin`) and press **Enter**.
4. Type the password for this administrator account and press **Enter**.

The FortiGate unit displays a command prompt (its hostname followed by a #). You can now enter CLI commands.



If three incorrect login or password attempts occur in a row, you will be disconnected. If this occurs, wait one minute, then reconnect to attempt the login again.

CLI-only features

As you can see in the [Feature / Platform Matrix](#), the entry level models have a number of features that are only available using the CLI, rather than appearing in the GUI.

You can open the CLI console so that it automatically opens to the object you wish to configure. For example, to edit a firewall policy, right-click on the policy in the policy list (**Policy & Objects > IPv4 Policy**) and select **Edit in CLI**. The CLI console will appear, with the commands to access this part of the configuration added automatically.

Once you have access to the CLI, you can enter instructions for specific tasks that can be found throughout the FortiOS Handbook. Options are also available at the top of the CLI Console to **Clear console**, **Download**, and **Copy to clipboard**.

Refer to the [CLI Reference](#) for a list of the available commands.

Command syntax

When entering a command, the CLI console requires that you use valid syntax and conform to expected input constraints. It will reject invalid commands.

Fortinet documentation uses the conventions below to describe valid command syntax.

Terminology

Each command line consists of a command word that is usually followed by configuration data or other specific item that the command uses or affects.

To describe the function of each word in the command line, especially if that nature has changed between firmware versions, Fortinet uses terms with the following definitions.

Command syntax terminology

- **Command** — A word that begins the command line and indicates an action that the FortiGate should perform on a part of the configuration or host on the network, such as `config` or `execute`. Together with other words, such as fields or values, that end when you press the **Enter** key, it forms a command line. Exceptions include multiline command lines, which can be entered using an escape sequence.
Valid command lines must be unambiguous if abbreviated. Optional words or other command line permutations are indicated by syntax notation.
- **Sub-command** — A `config` sub-command that is available only when nested within the scope of another command. After entering a command, its applicable sub-commands are available to you until you exit the scope of the command, or until you descend an additional level into another sub-command. Indentation is used to indicate levels of nested commands.
Not all top-level commands have sub-commands. Available sub-commands vary by their containing scope.
- **Object** — A part of the configuration that contains tables and /or fields. Valid command lines must be specific enough to indicate an individual object.
- **Table** — A set of fields that is one of possibly multiple similar sets which each have a name or number, such as an administrator account, policy, or network interface. These named or numbered sets are sometimes referenced by other parts of the configuration that use them.
- **Field** — The name of a setting, such as `ip` or `hostname`. Fields in some tables must be configured with values. Failure to configure a required field will result in an invalid object configuration error message, and the FortiGate will discard the invalid table.
- **Value** — A number, letter, IP address, or other type of input that is usually your configuration setting held by a field. Some commands, however, require multiple input values which may not be named but are simply entered in sequential order in the same command line. Valid input types are indicated by constraint notation.
- **Option** — A kind of value that must be one or more words from of a fixed set of options.

Indentation

Indentation indicates levels of nested commands, which indicate what other sub-commands are available from within the scope. The “`next`” and “`end`” lines are used to maintain a hierarchy and flow to CLI commands, especially helping to distinguish those commands with extensive sub-commands.

The “`next`” line is entered at the same indentation-level as the previous “`edit`”, to mark where you would like to finish that table entry and move on to the next table entry; doing so will not mean that you have “left” that sub-command.

next

Below is an example command, with a sub-command of `entries`:

```
config dlp filepattern
edit <1>
    set name <name>
    set comment [comment]
    config entries
        edit <2>
            set filter-type {pattern | type}
        next
    ←
```

After entering settings for <2> and entering `next`, the <2> table entry has been saved, and you be set back one level of indentation so you can continue to create more `entries` (if you wish).

This hierarchy is best indicated in the CLI console, as the example below is what displays in the console after entering `next`:

```
FGT60E1Q23456789 (entries) #
```



To go-back up an indentation-level from this point on (i.e. to finish configuring the `entries` sub-command), you **cannot** enter `next`; you must enter `end`.

end

Below is the same command and sub-command, except `end` has been entered instead of `next` after the sub-command:

```
config dlp filepattern
edit <1>
    set name <name>
    set comment [comment]
    config entries
        edit <2>
            set filter-type {pattern | type}
        end
    ←
```

Entering `end` will save the <2> table entry, but bring you out of the sub-command entirely; in this example, you would enter this when you don't wish to continue creating new `entries`.

Again, your hierarchy is best indicated by the CLI console. Below is what displays in the console after entering `end`:

```
FGT60E1Q23456789 (1) #
```

Notation

Brackets, braces, and pipes are used to denote valid permutations of the syntax. Constraint notations, such as `<address_ipv4>`, indicate which data types or string patterns are acceptable value input.

All syntax uses the following conventions:

Convention	Description
Square brackets []	<p>An optional word or series of words. For example:</p> <pre>[verbose {1 2 3}]</pre> <p>indicates that you may either omit or type both the word <code>verbose</code> and its accompanying option/s, such as <code>verbose 3</code>.</p> <p>See Optional values and ranges below for more information.</p>
Curly braces { }	<p>A word or series of words that is constrained to a set of options delimited by either vertical bars or spaces. You must enter at least one of the options, unless the set of options is surrounded by square brackets [].</p>
Mutually exclusive options - delimited by vertical bars	<p>Both mutually and non-mutually exclusive commands will use curly braces, as they provide multiple options, however mutually exclusive commands will divide each option with a pipe. This indicates that you are permitted to enter one option or the other:</p> <pre>{enable disable}</pre>
Non-mutually exclusive options - delimited by spaces	<p>Non-mutually exclusive commands do not use pipes to divide their options. In those circumstances, multiple options can be entered at once, as long as they are entered with a space separating each option:</p> <pre>{http https ping snmp ssh telnet}</pre>

Convention	Description
Angle brackets < >	<p>A word constrained by data type. The angled brackets contain a descriptive name followed by an underscore (_) and suffix that indicates the valid data type. For example, <retries_int>, indicates that you should enter a number of retries as an integer.</p> <p>Data types include:</p> <ul style="list-style-type: none"> • <xxx_name>: A name referring to another part of the configuration, such as <code>policy_A</code>. • <xxx_index>: An index number referring to another part of the configuration, such as 0 for the first static route. • <xxx_pattern>: A regular expression or word with wild cards that matches possible variations, such as <code>*@example.com</code> to match all email addresses ending in <code>@example.com</code>. • <xxx_fqdn>: A fully qualified domain name (FQDN), such as <code>mail.example.com</code>. • <xxx_email>: An email address, such as <code>admin@example.com</code>. • <xxx_ipv4>: An IPv4 address, such as <code>192.168.1.99</code>. • <xxx_v4mask>: A dotted decimal IPv4 netmask, such as <code>255.255.255.0</code>. • <xxx_ipv4mask>: A dotted decimal IPv4 address and netmask separated by a space, such as <code>192.168.1.99 255.255.255.0</code>. • <xxx_ipv4/mask>: A dotted decimal IPv4 address and CIDR-notation netmask separated by a slash, such as <code>192.168.1.1/24</code> • <xxx_ipv4range> : A hyphen (-)-delimited inclusive range of IPv4 addresses, such as <code>192.168.1.1-192.168.1.255</code>. • <xxx_ipv6>: A colon (:)-delimited hexadecimal IPv6 address, such as <code>3f2e:6a8b:78a3:0d82:1725:6a2f:0370:6234</code>. • <xxx_v6mask>: An IPv6 netmask, such as <code>/96</code>. • <xxx_ipv6mask>: A dotted decimal IPv6 address and netmask separated by a space. • <xxx_str>: A string of characters that is not another data type, such as <code>P@ssw0rd</code>. Strings containing spaces or special characters must be surrounded in quotes or use escape sequences. • <xxx_int>: An integer number that represents a metric, <code>minutes_int</code> for the number of minutes.

Optional values and ranges

Any field that is optional will use square-brackets, such as `set comment`. This is because it doesn't matter whether it's set or not. The overall config command will still successfully be taken.

Another example of where square-brackets would be used is to show that multiple options can be set, even intermixed with ranges. The example below shows a field that can be set to either a specific value or range, or multiple instances:

```
config firewall service custom
  set iprange <range1> [<range2> <range3> ...]
```

end

Sub-commands

Each command line consists of a command word that is usually followed by configuration data or other specific item that the command uses or affects:

```
get system admin
```

Sub-commands are available from within the scope of some commands. When you enter a sub-command level, the command prompt changes to indicate the name of the current command scope. For example, after entering:

```
config system admin
```

the command prompt becomes:

```
(admin) #
```

Applicable sub-commands are available to you until you exit the scope of the command, or until you descend an additional level into another sub-command.

For example, the `edit` sub-command is available only within a command that affects tables; the `next` sub-command is available only from within the `edit` sub-command:

```
config system interface
  edit port1
    set status up
  next
end
```

Sub-command scope is indicated by indentation.

Available sub-commands vary by command. From a command prompt within `config`, two types of sub-commands might become available:

- commands affecting fields
- commands affecting tables

Commands for tables

clone <table>

Clone (or make a copy of) a table from the current object.

For example, in `config firewall policy`, you could enter the following command to clone security policy 27 to create security policy 30:

```
clone 27 to 30
```

In `config antivirus profile`, you could enter the following command to clone an antivirus profile named `av_pro_1` to create a new antivirus profile named `av_pro_2`:

```
clone av_pro_1 to av_pro_2
```

`clone` may not be available for all tables.

delete <table>

Remove a table from the current object.

For example, in `config system admin`, you could delete an administrator account named `newadmin` by typing `delete newadmin` and pressing Enter. This deletes `newadmin` and all its fields, such as `newadmin's first-name` and `email-address`.

`delete` is only available within objects containing tables.

edit <table>

Create or edit a table in the current object.

For example, in `config system admin`:

- edit the settings for the default `admin` administrator account by typing `edit admin`.
- add a new administrator account with the name `newadmin` and edit `newadmin's` settings by typing `edit newadmin`.

`edit` is an interactive sub-command: further sub-commands are available from within `edit`.

`edit` changes the prompt to reflect the table you are currently editing.

`edit` is only available within objects containing tables.

In objects such as security policies, `<table>` is a sequence number. To create a new entry without the risk of overwriting an existing one, enter `edit 0`. The CLI initially confirms the creation of entry 0, but assigns the next unused number after you finish editing and enter `end`.

end

Save the changes to the current object and exit the `config` command. This returns you to the top-level command prompt.

get

List the configuration of the current object or table.

- In objects, `get` lists the table names (if present), or fields and their values.
- In a table, `get` lists the fields and their values.

For more information on `get` commands, see the [CLI Reference](#).

purge

Remove all tables in the current object.

For example, in `config user local`, you could type `get` to see the list of user names, then type `purge` and then `y` to confirm that you want to delete all users.

`purge` is only available for objects containing tables.

Caution: Back up the FortiGate before performing a `purge`. `purge` cannot be undone. To restore purged tables, the configuration must be restored from a backup.

Caution: Do not purge `system interface` or `system admin` tables. `purge` does not provide default tables. This can result in being unable to connect or log in, requiring the FortiGate to be formatted and restored.

rename <table> to <table>

Rename a table.

For example, in `config system admin`, you could rename `admin3` to `fwadmin` by typing `rename admin3 to fwadmin`.

`rename` is only available within objects containing tables.

show

Display changes to the default configuration. Changes are listed in the form of configuration commands.

Example of table commands

From within the `system admin` object, you might enter:

```
edit admin_1
```

The CLI acknowledges the new table, and changes the command prompt to show that you are now within the `admin_1` table:

```
new entry 'admin_1' added
(admin_1) #
```

Commands for fields**abort**

Exit both the `edit` and/or `config` commands without saving the fields.

append

Add an option to an existing list.

end

Save the changes made to the current table or object fields, and exit the `config` command (to exit without saving, use `abort` instead).

get

List the configuration of the current object or table.

- In objects, `get` lists the table names (if present), or fields and their values.
- In a table, `get` lists the fields and their values.

move	Move an object within a list, when list order is important. For example, rearranging security policies within the policy list.
next	<p>Save the changes you have made in the current table's fields, and exit the <code>edit</code> command to the object prompt (to save and exit completely to the root prompt, use <code>end</code> instead).</p> <p><code>next</code> is useful when you want to create or edit several tables in the same object, without leaving and re-entering the <code>config</code> command each time.</p> <p><code>next</code> is only available from a table prompt; it is not available from an object prompt.</p>
select	<p>Clear all options except for those specified.</p> <p>For example, if a group contains members A, B, C, and D and you remove all users except for B, use the command <code>select member B</code>.</p>
set <field> <value>	<p>Set a field's value.</p> <p>For example, in <code>config system admin</code>, after typing <code>edit admin</code>, you could type <code>set password newpass</code> to change the password of the <code>admin</code> administrator to <code>newpass</code>.</p> <p>Note: When using <code>set</code> to change a field containing a space-delimited list, type the whole new list. For example, <code>set <field> <new-value></code> will replace the list with the <code><new-value></code> rather than appending <code><new-value></code> to the list.</p>
show	Display changes to the default configuration. Changes are listed in the form of configuration commands.
unselect	Remove an option from an existing list.
unset <field>	<p>Reset the table or object's fields to default values.</p> <p>For example, in <code>config system admin</code>, after typing <code>edit admin</code>, typing <code>unset password</code> resets the password of the <code>admin</code> administrator account to the default (in this case, no password).</p>

Example of field commands

To assign the value `my1stExamplePassword` to the `password` field, enter the following command from within the `admin_1` table:

```
set password my1stExamplePassword
```

Next, to save the changes and edit the next administrator's table, enter the `next` command.

Permissions

Access profiles control which CLI commands an administrator account can access. Access profiles assign either read, write, or no access to each area of FortiOS. To view configurations, you must have read access. To make changes, you must have write access. So, depending on the account used to log in to the FortiGate unit, you may not have complete access to all CLI commands. For complete access to all commands, you must log in with the administrator account named `admin`.

Unlike other administrator accounts, the `admin` account exists by default and cannot be deleted. The `admin` account is similar to a root administrator account that always has full permission to view and change all FortiGate configuration options, including viewing and changing all other administrator accounts. Its name and permissions cannot be changed. It is the only administrator account that can reset another administrator's password without being required to enter that administrator's existing password.



Set a strong password for the `admin` account, and change the password regularly. By default, this administrator account has no password. Failure to maintain the password of the `admin` account could compromise the security of your FortiGate.

Tips

Basic features and characteristics of the CLI environment provide support and ease of use for many CLI tasks.

Help

To display brief help during command entry, press the question mark (?) key.

- Press the question mark (?) key at the command prompt to display a list of the commands available and a description of each command.
- Type a word or part of a word, then press the question mark (?) key to display a list of valid word completions or subsequent words, and to display a description of each.

Shortcuts and key commands

Keys	Action
?	List valid word completions or subsequent words. If multiple words could complete your entry, display all possible completions with helpful descriptions of each.
Tab	Complete the word with the next available match. Press the Tab key multiple times to cycle through available matches.
Up arrow, or Ctrl + P	Recall the previous command. Command memory is limited to the current session.
Down arrow, or Ctrl + N	Recall the next command.
Left or Right arrow	Move the cursor left or right within the command line.
Ctrl + A	Move the cursor to the beginning of the command line.
Ctrl + E	Move the cursor to the end of the command line.
Ctrl + B	Move the cursor backwards one word.
Ctrl + F	Move the cursor forwards one word.
Ctrl + D	Delete the current character.

Keys	Action
Ctrl + C	<p>Abort current interactive commands, such as when entering multiple lines.</p> <p>If you are not currently within an interactive command such as <code>config</code> or <code>edit</code>, this closes the CLI connection.</p>
\ then Enter	<p>Continue typing a command on the next line for a multiline command.</p> <p>For each line that you want to continue, terminate it with a backslash (\). To complete the command line, terminate it by pressing the spacebar and then the Enter key, without an immediately preceding backslash.</p>

Command abbreviation

You can abbreviate words in the command line to their smallest number of non-ambiguous characters.

For example, the command `get system status` could be abbreviated to `g sy stat`.

Adding and removing options from lists

When adding options to a list, such as a user group, using the `set` command will remove the previous configuration. For example, if you wish to add user D to a user group that already contains members A, B, and C, the command would need to be `set member A B C D`. If only `set member D` was used, then all former members would be removed from the group.

However, there are additional commands which can be used instead of `set` for changing options in a list.

Additional commands for lists

append	<p>Add an option to an existing list.</p> <p>For example, <code>append member</code> would add user D to a user group while all previous group members are retained</p>
select	<p>Clear all options except for those specified.</p> <p>For example, if a group contains members A, B, C, and D and you remove all users except for B, use the command <code>select member B</code>.</p>
unselect	<p>Remove an option from an existing list.</p> <p>For example, <code>unselect member A</code> would remove member A from a group while all previous group members are retained.</p>

Environment variables

The CLI supports the following environment variables. Variable names are case-sensitive.

Environment variables

\$USERFROM	The management access type (<code>ssh</code> , <code>telnet</code> , <code>jsconsole</code> for the CLI Console widget in the web-based manager, and so on) and the IP address of the administrator that configured the item.
\$USERNAME	The account name of the administrator that configured the item.
\$SerialNum	The serial number of the FortiGate unit.

For example, the FortiGate unit's host name can be set to its serial number:

```
config system global
    set hostname $SerialNum
end
```

Special characters

The following special characters, also known as reserved characters, are not permitted in most CLI fields:

< > () # ' “

You may be able to enter special characters as part of a string's value by using a special command, enclosing it in quotes, or preceding it with an escape sequence — in this case, a backslash (\) character.

In other cases, different keystrokes are required to input a special character. If you need to enter **?** as part of config, you first need to input **CTRL-V**. If you enter **?** without first using **CTRL-V**, the question mark has a different meaning in the CLI; it will show available command options in that section.

For example, if you enter **?** without **CTRL-V**:

```
edit "*.xe
token line: Unmatched double quote.
```

If you enter **?** with **CTRL-V**:

```
edit "*.xe?"
new entry '*.xe?' added
```

Entering special characters

Character	Keys
?	Ctrl + V then ?
Tab	Ctrl + V then Tab

Character	Keys
Space (to be interpreted as part of a string value, not to end the string)	Enclose the string in quotation marks: "Security Administrator". Enclose the string in single quotes: 'Security Administrator'. Precede the space with a backslash: Security\ Administrator.
'	\'
(to be interpreted as part of a string value, not to end the string)	
"	\"
(to be interpreted as part of a string value, not to end the string)	
\	\\

Using grep to filter get and show command output

In many cases, the `get` and `show` (and `diagnose`) commands may produce a large amount of output. If you are looking for specific information in a large `get` or `show` command output, you can use the `grep` command to filter the output to only display what you are looking for. The `grep` command is based on the standard UNIX `grep`, used for searching text output based on regular expressions.

Use the following command to display the MAC address of the FortiGate unit internal interface:

```
get hardware nic internal | grep Current_HWaddr
Current_HWaddr           00:09:0f:cb:c2:75
```

Use the following command to display all TCP sessions in the session list and include the session list line number in the output:

```
get system session list | grep -n tcp
```

Use the following command to display all lines in HTTP replacement message commands that contain URL (upper or lower case):

```
show system replacemsg http | grep -i url
```

There are three additional options that can be applied to `grep`:

```
-A <num> After
-B <num> Before
-C <num> Context
```

The option `-f` is also available to support contextual output, in order to show the complete configuration. The following example shows the difference in output when `-f` option is used versus when it is not.

Using -f:

```

show | grep -f ldap-group1
config user group
  edit "ldap-group1"
    set member "pc40-LDAP"
  next
end
config firewall policy
  edit 2
    set srcintf "port31"
    set dstintf "port32"
    set srcaddr "all"
    set action accept
    set identity-based enable
    set nat enable
    config identity-based-policy
      edit 1
        set schedule "always"
        set groups "ldap-group1"
        set dstaddr "all"
        set service "ALL"
      next
    end
  next
end

```

Without using -f:

```

show | grep ldap-group1
edit "ldap-group1"
  set groups "ldap-group1"

```

Language support and regular expressions

Characters such as ñ, é, symbols, and ideographs are sometimes acceptable input. Support varies by the nature of the item being configured. CLI commands, objects, field names, and options must use their exact ASCII characters, but some items with arbitrary names or values may be input using your language of choice. To use other languages in those cases, you must use the correct encoding.

Input is stored using Unicode UTF-8 encoding but is not normalized from other encodings into UTF-8 before it is stored. If your input method encodes some characters differently than in UTF-8, your configured items may not display or operate as expected.

Regular expressions are especially impacted. Matching uses the UTF-8 character values. If you enter a regular expression using another encoding, or if an HTTP client sends a request in an encoding other than UTF-8, matches may not be what you expect.

For example, with Shift-JIS, backslashes (\) could be inadvertently interpreted as the symbol for the Japanese yen (¥) and vice versa. A regular expression intended to match HTTP requests containing money values with a yen symbol therefore may not work if the symbol is entered using the wrong encoding.

For best results, you should:

- use UTF-8 encoding, or
- use only the characters whose numerically encoded values are the same in UTF-8, such as the US-ASCII characters that are also encoded using the same values in ISO 8859-1, Windows code page 1252, Shift-JIS and other encodings, or
- for regular expressions that must match HTTP requests, use the same encoding as your HTTP clients.



HTTP clients may send requests in encodings other than UTF-8. Encodings usually vary by the client's operating system or input language. If you cannot predict the client's encoding, you may only be able to match any parts of the request that are in English, because regardless of the encoding, the values for English characters tend to be encoded identically. For example, English words may be legible regardless of interpreting a web page as either ISO 8859-1 or as GB2312, whereas simplified Chinese characters might only be legible if the page is interpreted as GB2312.

If you configure your FortiGate unit using other encodings, you may need to switch language settings on your management computer, including for your web browser or Telnet/SSH client. For instructions on how to configure your management computer's operating system language, locale, or input method, see its documentation.

If you choose to configure parts of the FortiGate unit using non-ASCII characters, verify that all systems interacting with the FortiGate unit also support the same encodings. You should also use the same encoding throughout the configuration if possible in order to avoid needing to switch the language settings of the web-based manager and your web browser or Telnet/SSH client while you work.

Similarly to input, your web browser or CLI client should normally interpret display output as encoded using UTF-8. If it does not, your configured items may not display correctly in the GUI or CLI. Exceptions include items such as regular expressions that you may have configured using other encodings in order to match the encoding of HTTP requests that the FortiGate unit receives.

To enter non-ASCII characters in the CLI console:

1. On your management computer, start your web browser and go to the URL for the FortiGate unit's GUI.
2. Configure your web browser to interpret the page as UTF-8 encoded.
3. Log in to the FortiGate unit.
4. Open the CLI Console from the upper right-hand corner.
5. In the title bar of the **CLI Console** widget, click **Edit** (the pencil icon).
6. Enable **Use external command input box** and select **OK**.
7. The **Command** field appears below the usual input and display area of the **CLI Console**.
8. Type a command in this field and press **Enter**.

In the display area, the **CLI Console** widget displays your previous command interpreted into its character code equivalent, such as:

```
edit \743\601\613\743\601\652
```

and the command's output.

To enter non-ASCII characters in a Telnet/SSH client

1. On your management computer, start your Telnet or SSH client.
2. Configure your Telnet or SSH client to send and receive characters using UTF-8 encoding.
Support for sending and receiving international characters varies by each Telnet/SSH client. Consult the documentation for your Telnet/SSH client.

3. Log in to the FortiGate unit.
4. At the command prompt, type your command and press **Enter**.

You may need to surround words that use encoded characters with single quotes (').

Depending on your Telnet/SSH client's support for your language's input methods and for sending international characters, you may need to interpret them into character codes before pressing Enter.

For example, you might need to enter:

```
edit '\743\601\613\743\601\652'
```

5. The CLI displays your previous command and its output.

Screen paging

You can configure the CLI to pause after displaying each page's worth of text when displaying multiple pages of output. When the display pauses, the last line displays `--More--`. You can then either:

- press the **spacebar** to display the next page.
- type **q** to truncate the output and return to the command prompt.

This may be useful when displaying lengthy output, such as the list of possible matching commands for command completion, or a long list of settings. Rather than scrolling through or possibly exceeding the buffer of your terminal emulator, you can simply display one page at a time.

To configure the CLI Console to pause display when the screen is full:

```
config system console
  set output more
end
```

Baud rate

You can change the default baud rate of the local console connection.

To change the baud rate enter the following commands:

```
config system console
  set baudrate {9600 | 19200 | 38400 | 57600 | 115200}
end
```

Editing the configuration file on an external host

You can edit the FortiGate configuration on an external host by first backing up the configuration file to a TFTP server. Then edit the configuration file and restore it to the FortiGate unit.

Editing the configuration on an external host can be timesaving if you have many changes to make, especially if your plain text editor provides advanced features such as batch changes.

To edit the configuration on your computer:

1. Use `execute backup` to download the configuration file to a TFTP server, such as your management computer.
2. Edit the configuration file using a plain text editor that supports Unix-style line endings.



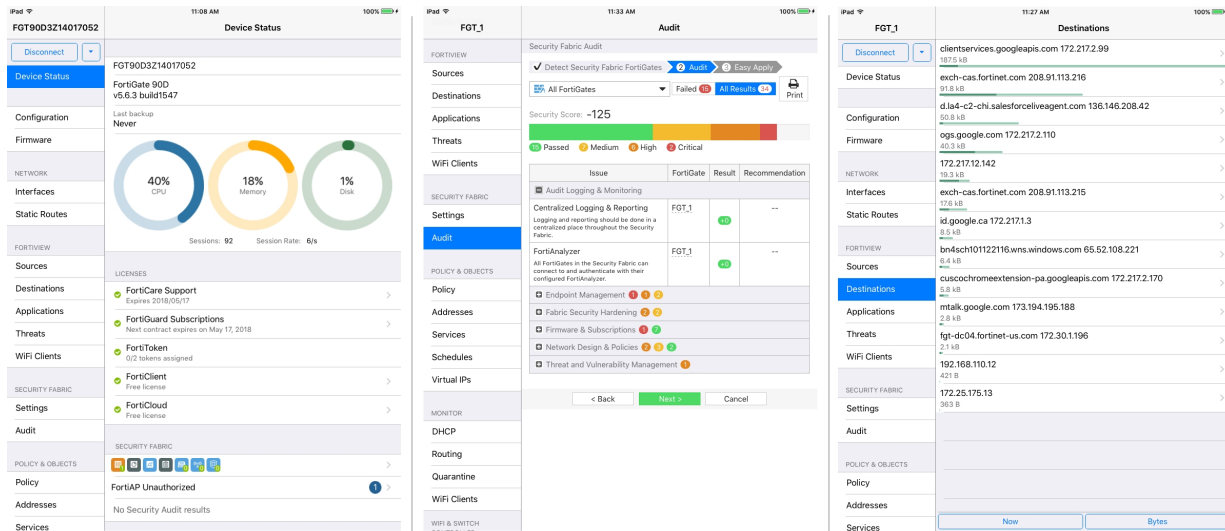
Do not edit the first line. The first line(s) of the configuration file (preceded by a # character) contains information about the firmware version and FortiGate model. If you change the model number, the FortiGate unit will reject the configuration file when you attempt to restore it.

3. Use `execute restore` to upload the modified configuration file back to your FortiGate.

The FortiGate downloads the configuration file and checks that the model information is correct. If it is correct, the FortiGate unit loads the configuration file and checks each command for errors. If a command is invalid, the FortiGate unit ignores the command. If the configuration file is valid, the FortiGate unit restarts and loads the new configuration.

FortiExplorer for iOS

FortiExplorer for iOS is a user-friendly application that helps you to quickly and easily configure, manage, and monitor FortiGate appliances using an iOS device. FortiExplorer lets you rapidly provision, deploy, and monitor Security Fabric components including FortiGate, FortiWiFi, and FortiAP devices.



FortiExplorer for iOS requires iOS 9.3 or later and is compatible with iPhone, iPad, and iPod Touch. It is supported by FortiOS 5.6+ and is [only available on the App Store](#) for iOS devices.

Advanced features are available with the purchase of FortiExplorer Pro. Paid features include the ability to add more than two devices and the ability to download firmware images from FortiCare.

Up to six members can use this app with 'Family Sharing' enabled in the App Store.

Getting started with FortiExplorer

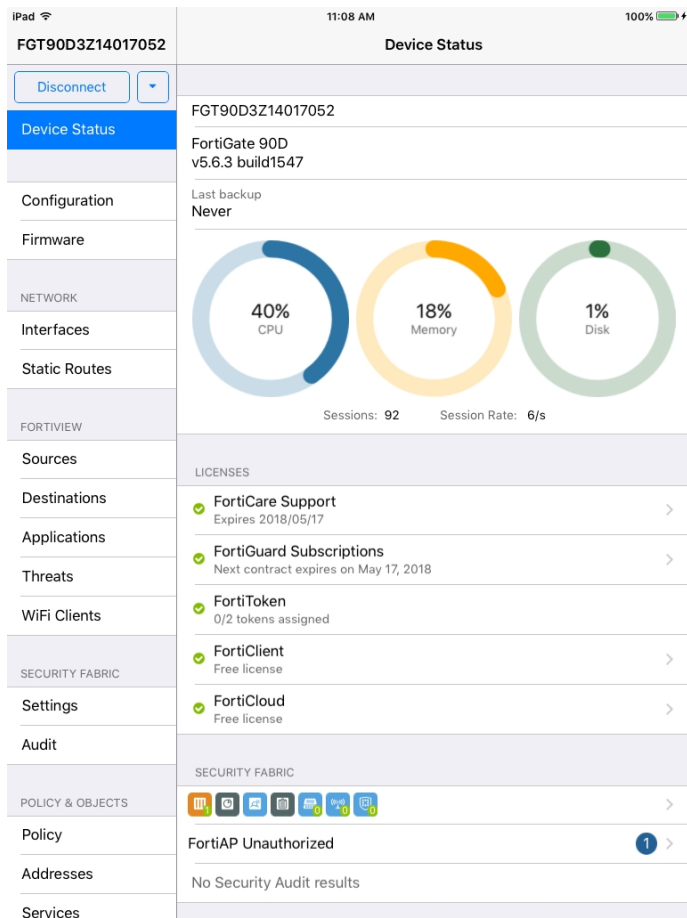
If your FortiGate is accessible on the wireless network, you can connect to it using FortiExplorer provided that your iOS device is on the same network (see [Connecting FortiExplorer to a FortiGate via WiFi](#)). Otherwise, you will need to physically connect your iOS device to the FortiGate using a USB cable (see below).

Connecting FortiExplorer to a FortiGate via USB


For the purpose of this document, we assume that you are just getting started; you do not have access to the FortiGate over the wireless network, and the FortiGate is in its factory configuration.

1. Connect your iOS device to your FortiGate's USB management port.
If prompted on your iOS device, **Trust** this 'computer'.
2. Open the FortiExplorer app and select your FortiGate from the list under **USB Attached Device**.
3. On the **Login** screen, select **USB**.
4. Enter the default **Username** (`admin`) and leave the **Password** field blank.
5. You can opt to **Remember Password**. Tap **Done** when you are ready.

FortiExplorer opens the FortiGate management interface to the **Device Status** page:



6. Go to **Network > Interfaces** and configure the WAN interface(s).
 In the example, the **wan1** interface **Address** mode is set to **DHCP** by default. Set it to **Manual** and enter its **Address**, **Netmask**, and **Default Gateway**, and then **Apply** your changes.



7. (Optional) Configure **Administrative Access** to allow **HTTP** and **HTTPS** access.
 This will allow administrators to access the FortiGate web-based manager using a web browser.

8. Go to **Network > Interfaces** and configure the local network (internal) interface. Set the **Address** mode as before and configure **Administrative Access** if desired.
9. Configure a **DHCP Server** for the internal network subnet.

Return to the internal interface using the button at the top of the screen.

10. Go to **Network > Static Routes** and configure the static route to the gateway.

11. Go to **Policy & Objects > Policy** and edit the Internet access policy. As a best practice, provide a **Name** for the policy, enable the desired **Security Profiles**, and configure **Logging Options**. Select **OK** to finalize.

iPad 11:11 AM 100%

FGT90D3Z14017052 Policy

Disconnect

Device Status

Configuration

Firmware

NETWORK

Interfaces

Static Routes

FORTIVIEW

Sources

Destinations

Applications

Threats

WiFi Clients

SECURITY FABRIC

Settings

Audit

POLICY & OBJECTS

Policy

Addresses

Services

Edit Policy

Name

Incoming Interface internal

Outgoing Interface wan1

Source all

Destination all

Schedule always

Service ALL

Action ACCEPT DENY LEARN

Firewall / Network Options

NAT

IP Pool Configuration Use Outgoing Interface Address Use Dynamic IP Pool

Security Profiles

AntiVirus

Web Filter

DNS Filter

Application Control

SSL Inspection

Logging Options

Log Allowed Traffic Security Events All Sessions

Capture Packets

Comments Write a comment... 0/1023

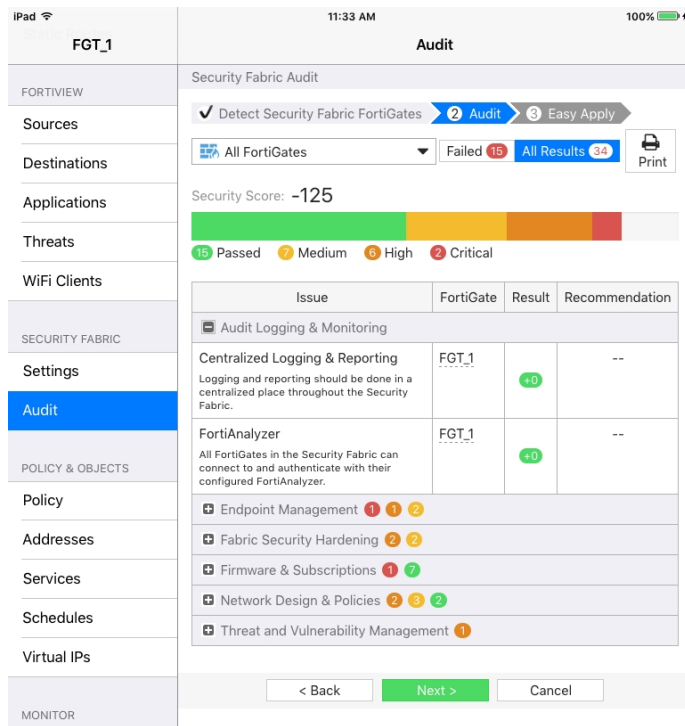
Enable this policy

OK Cancel

Running a Security Fabric Rating

The FortiGate is now configured in a very basic state. Once you've configured the other potential elements of your network, such as other **Interfaces**, **Schedules**, or **Managed FortiAPs**, it is recommended that you run a **Security Fabric Rating** to identify potential vulnerabilities and highlight best practices that could be used to improve your network's overall security and performance.

Go to **Security Fabric > Security Rating** and follow the steps to determine a **Security Score** for the selected device(s). The results should identify issues ranging from Medium to Critical importance, and may provide recommended actions where possible.



Connecting FortiExplorer to a FortiGate via WiFi

If your FortiGate is accessible on the wireless network, you can connect to it using FortiExplorer provided that your iOS device is on the same network. Assuming this is the case:

1. Open the FortiExplorer app and select **Add** from the **Devices** page.
2. Enter the **Host** information and appropriate **Username** and **Password** credentials. If necessary, change the default **Port** number, and opt to **Remember Password**.

The screenshot shows the FortiExplorer app login screen. At the top, there are buttons for 'Cancel', 'Login', and 'Done'. Below these are input fields for 'Host' (172.25.176.56), 'Port' (443), 'Username' (admin), and 'Password'. There is also a 'Remember Password' toggle switch.

3. If the FortiGate device identity cannot be verified, click **Connect** at the prompt. FortiExplorer opens the FortiGate management interface to the **Device Status** page.

Upgrading to FortiExplorer Pro

Paid features provided with the purchase of FortiExplorer Pro include the ability to add more than two devices and the ability to download firmware images from FortiCare.

- To upgrade to FortiExplorer Pro, open the FortiExplorer app, go to **Settings** and select **Upgrade to FortiExplorer Pro**. Follow the on-screen prompts.

LED specifications

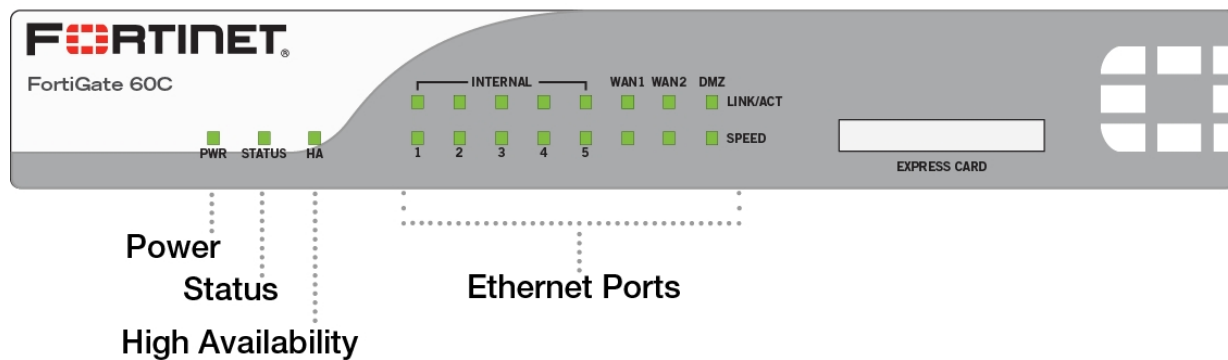
The following section includes information regarding FortiGate LED status indicators.

- [Sample FortiGate faceplates](#)
- [LED status codes](#)
- [About alarm levels](#)
- [LED status codes for ports](#)


Sample FortiGate faceplates

The faceplates indicate where the LEDs are typically found on desktop and mid-range FortiGate models.

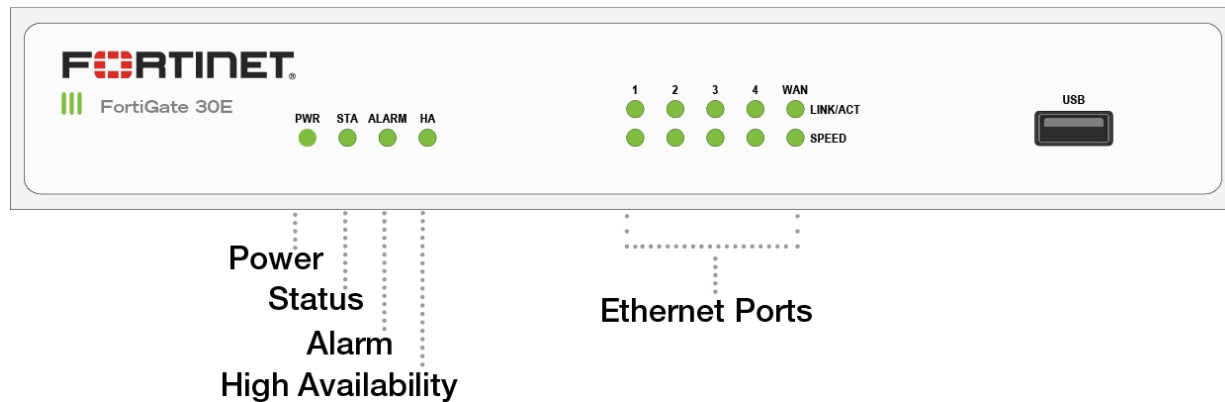
FortiGate 60C



FortiGate 100D



FortiGate 30E



LED status codes

LABEL	STATE	MEANING
PWR	Green	Power is On.
	Off	Power is Off.
STA	Green	Normal status.
	Flashing Green	Booting Up. If the FortiGate has a reset button, Flashing Green also means that the reset button was used.
	Red	The FortiGate has a critical alarm (see About Alarm Levels).
ALARM	Off	No alarms or the FortiGate has a minor alarm.
	Amber	The FortiGate has a major alarm.
	Red	The FortiGate has a critical alarm. The status LED will also be red.
		More information at About Alarm Levels .
HA	Green	FortiGate is operating in an FGCP HA cluster.
	Red	A failover has occurred.
	Off	HA not configured.
		Failover operation feature not available in all units.

LABEL	STATE	MEANING
WIFI	Green	Wireless port is active.
	Flashing Green	Wireless interface is transmitting and receiving data.
	Off	Wireless interface is down.

About alarm levels

Minor, major, and critical alarms are defined based on IPMI, ATCA, and Telco standards for naming alarms.

- A minor alarm (also called an IPMI non-critical (NC) alarm) indicates a temperature or a power level outside of the normal operating range that is not considered a problem. In the case of a minor temperature alarm, the system could respond by increasing fan speed. A non-critical threshold can be an upper non-critical (UNC) threshold (for example, a high temperature or a high power level) or a lower non-critical (LNC) threshold (for example, a low power level). The LEDs do not indicate minor alarms since user intervention is not required.
- A major alarm (also called an IPMI critical or critical recoverable (CR) alarm) indicates that the system itself cannot correct the cause for the alarm and that intervention is required. For example, the cooling system cannot provide enough cooling to reduce the temperature. It could also mean that conditions (e.g. temperature) are approaching the outside limit of the allowed operating range. A critical threshold can also be an upper critical (UC) threshold (e.g. a high temperature or a high power level) or a lower critical (LC) threshold (e.g. a low power level).
- A critical alarm (also called an IPMI non-recoverable (NR) alarm) indicates detection of a temperature or power level that is outside of the allowed operating range and could potentially cause physical damage.

LED status codes for ports

TYPE OF PORT	STATE	MEANING
Ethernet Ports Link / Activity	Green	Connected.
	Flashing Green	Transmitting and receiving data.
	Off	No link established.
		On FortiGate models with front-facing ports, this LED is to the left of the port. On FortiGate models with ports at the back of the device, this LED is in the upper row.
SFP Ports	Green	Connected.
	Flashing Green	Transmitting and receiving data.
	Off	No link established.

TYPE OF PORT	STATE	MEANING
Ethernet Ports Speed	Green	Connected at 1Gbps.
	Amber	Connected at 100Mbps.
	Off	Not connected or connected at 10Mbps.
		On FortiGate models with front-facing ports, this LED is to the right of the port. On FortiGate models with ports at the back of the device, this LED is in the lower row.

Inspection mode

To control your FortiGate's security profile inspection mode in FortiOS 6.0, you can select **Flow-based** or **Proxy** inspection modes from **System > Settings**. Having control over flow and proxy mode is helpful if you want to ensure that only flow inspection mode is used.

In most cases proxy mode is preferred because more security profile features are available along with more configuration options for these individual features. Some implementations, however, may require all security profile scanning to only use flow mode. In this case, you can set your FortiGate to flow mode knowing that proxy mode inspection will not be used.

Setting up the FortiGate to operate in these new modes (or to operate in the other available operating modes) involves going to **System > Settings** and changing the **Inspection Mode** and **NGFW Mode**.

NGFW mode simplifies applying application control and web filtering to traffic by allowing you to add applications and web filtering profiles directly to policies.

Transparent proxy allows you to apply web authentication to HTTP traffic without using the explicit proxy.

Changing inspection and policy modes

To change inspection modes, go to **System > Settings**. You can select **Flow-based** or **Proxy** inspection modes.

NGFW mode

When you select **Flow-based** as the **Inspection Mode**, you have the option to select an **NGFW Mode**. In **NGFW Profile-based** mode, you configure Application Control and Web-Filtering profiles in **Security Profiles** and then apply them to a policy.

In **Policy-based** mode, you add applications and web filtering profiles directly to a policy without having to first create and configure Application Control or Web Filtering profiles.

When you change to **Flow-based** inspection, all proxy mode profiles are converted to flow mode, and proxy settings are removed. In addition, proxy-mode only features (for example, Web Application Profile) are removed from the GUI.

If your FortiGate has multiple VDOMs, you can set the inspection mode independently for each VDOM. Go to **System > VDOM**. Click **Edit** for the VDOM you want to change and select the **Inspection Mode**.

CLI syntax

You can use the following CLI command to configure NGFW mode:

```
config system settings
  set inspection-mode flow
  set ngfw-mode {profile-based | policy-based}
  set ssl-ssh-profile "certificate-inspection"
end
```

Security profile features mapped to inspection mode

The table below lists FortiOS security profile features and shows whether they are available in flow-based or proxy-based inspection modes.

Security Profile Feature	Flow-based inspection	Proxy-based inspection
AntiVirus	x	x
Web Filter	x	x
DNS Filter	x	x
Application Control	x	x
Intrusion Protection	x	x
Anti-Spam		x
Data Leak Protection		x
VoIP		x
ICAP		x
Web Application Firewall		x
FortiClient Profiles	x	x
Proxy Options	x	x
SSL Inspection	x	x
SSH Inspection		x
Web Rating Overrides	x	x
Web Profile Overrides		x

From the GUI, you can only configure antivirus and web filter security profiles in proxy mode. From the CLI, you can configure flow-based antivirus profiles, web filter profiles, and DLP profiles and they will appear on the GUI and include their inspection mode setting. Flow-based profiles created when in flow mode are still available when you switch to proxy mode.

In flow mode, antivirus and web filter profiles only include flow-mode features. Web filtering and virus scanning is still done with the same engines and to the same accuracy, but some inspection options are limited or not available in flow mode. Application control, intrusion protection, and FortiClient profiles are not affected when switching between flow and proxy mode.

Even though VoIP profiles are not available from the GUI in flow mode, the FortiGate can process VoIP traffic. In this case the appropriate session helper is used (for example, the SIP session helper).

Setting flow or proxy mode doesn't change the settings available from the CLI. However, when in flow mode you can't save security profiles that are set to proxy mode.

You can also add proxy-only security profiles to firewall policies from the CLI. So, for example, you can add a VoIP profile to a security policy that accepts VoIP traffic. This practice isn't recommended because the setting will not be visible from the GUI.

Proxy mode and flow mode antivirus and web filter profile options

The following tables list the antivirus and web filter profile options available in proxy and flow modes.

Antivirus features in proxy and flow mode

Feature	Proxy	Flow
Scan Mode (Quick or Full)		x
Detect viruses (Block or Monitor)	x	x
Inspected protocols	x	(all relevant protocols are inspected)
Inspection Options	x	x (not available for quick scan mode)
Treat Windows Executables in Email Attachments as Viruses	x	x
Send Files to FortiSandbox Appliance for Inspection	x	x
Use FortiSandbox Database	x	x
Include Mobile Malware Protection	x	x

Web filter features in proxy and flow mode

Feature	Proxy	Flow
FortiGuard category based filter	x	x (show, allow, monitor, block)
Category Usage Quota	x	
Allow users to override blocked categories (on some models)	x	
Search Engines	x	

Feature		Proxy	Flow
	Enforce 'Safe Search' on Google, Yahoo!, Bing, Yandex	x	
	Restrict YouTube Access	x	
	Log all search keywords	x	
Static URL Filter		x	x
	Block invalid URLs	x	
	URL Filter	x	x
	Block malicious URLs discovered by FortiSandbox	x	x
	Web Content Filter	x	x
Rating Options		x	x
	Allow websites when a rating error occurs	x	x
	Rate URLs by domain and IP Address	x	x
	Block HTTP redirects by rating	x	
	Rate images by URL	x	
Proxy Options		x	
	Restrict Google account usage to specific domains	x	
	Provide details for blocked HTTP 4xx and 5xx errors	x	
	HTTP POST Action	x	
	Remove Java Applets	x	
	Remove ActiveX	x	
	Remove Cookies	x	
	Filter Per-User Black/White List	x	

Basic administration

This section contains information about basic FortiGate administration that can be done after you have installed the unit in your network.

While this section mainly focuses on accomplishing tasks with the GUI, some tasks include instructions to use the CLI. You can access the CLI using the GUI or FortiExplorer, or via SSH or Telnet connection. For more information about the CLI, see [Using the CLI](#).

Registration

In order to have full access to Fortinet Support and FortiGuard Services, you must register your FortiGate.

Registering your FortiGate:

1. Go to the **Dashboard** and locate the **Licenses** widget.
2. Click on **FortiCare Support** to display a pop-up window and **Register**.
3. In the pop-up window, either use an existing Fortinet Support account or create a new one. Select your **Country** and **Reseller**.
4. Select **OK**.

FortiGate platforms do not impose any limitations on the number or type of customers, users, devices, IP addresses, or number of VPN clients being served by the platform. Such factors are limited solely by the hardware capacity of each given model.

System Settings

There are several system settings that should be configured once your FortiGate is installed:

- [Default administrator password](#)
- [Settings](#)
 - [Changing the host name](#)
 - [System Time](#)
 - [Administration Settings](#)
 - [Password Policy](#)
 - [View Settings](#)
- [Administrator password retries and lockout time](#)

Default administrator password

By default, your FortiGate has an administrator account set up with the username `admin` and no password. In order to prevent unauthorized access to the FortiGate, it is highly recommended that you add a password to this account.

To change the default password:

1. Go to **System > Administrators**.
2. Edit the **admin** account.

3. Select **Change Password**.
4. Enter the **New Password** and re-enter the password for confirmation.
5. Select **OK**.

For details on selecting a password and password best practices, see the section on [Passwords](#).

It is also recommended to change the user name of this account; however, since you cannot change the user name of an account that is currently in use, a second administrator account will need to be created in order to do this. For more information about creating and using administrator accounts, see the Administrators section of the [System Administration](#) chapter.

Settings

Settings can be accessed by going to **System > Settings**. On this page, you can change the **Host name**, set the system time and identify time zone in **System Time**, configure HTTP, HTTPS, SSH, and Telnet ports as well as idle timeout in **Administration Settings**, designate the **Password Policy**, and manage display options and designate inspection mode in **View Settings**.

Changing the host name

The host name of your FortiGate appears in the **Hostname** row in the **System Information** widget on the Dashboard. The host name also appears at the CLI prompt when you are logged in to the CLI, and as the SNMP system name.

To change the host name on the FortiGate

Go to **System > Settings** and type in the new name in the **Host name** row. The only administrators that can change a FortiGate's host name are administrators whose admin profiles permit system configuration write access. If the FortiGate is part of an HA cluster, you should use a unique host name to distinguish the FortiGate from others in the cluster.

System Time

For effective scheduling and logging, the FortiGate system time and date should be accurate. You can either manually set the system time and date or configure the FortiGate to automatically synchronize with a Network Time Protocol (NTP) server.

NTP enables you to keep the FortiGate time synchronized with other network systems. By enabling NTP on the FortiGate, FortiOS will check with the NTP server you select at the configured intervals. This will also ensure that logs and other time-sensitive settings on the FortiGate are correct.

The FortiGate maintains its internal clock using a built-in battery. At start up, the time reported by the FortiGate will indicate the hardware clock time, which may not be accurate. When using NTP, the system time might change after the FortiGate has successfully obtained the time from a configured NTP server.



By default, FortiOS has the daylight savings time configuration enabled. The system time must be manually adjusted after daylight saving time ends. To disable DST, enter the following commands in the CLI:

```
config system global
  set dst disable
end
```

To set the date and time

1. Go to the **System > Settings**.
2. Under **System Time**, select your **Time Zone** by using the drop-down menu.
3. **Set Time** by either selecting **Synchronize with NTP Server** or **Manual settings**. If you select synchronization, you can either use the default FortiGuard servers or specify a custom server. You can also set the **Sync interval**.
4. If you use an NTP server, you can identify a specific interface for this self-originating traffic by enabling **Setup device as local NTP server**.
5. Select **Apply**.

Administration Settings

In order to improve security, you can change the default port configurations for administrative connections to the FortiGate. When connecting to the FortiGate when the port has changed, the port must be included, such as `https://<ip_address>:<port>`. For example, if you are connecting to the FortiGate using port 99, the URL would be `https://192.168.1.99:99`.

To configure the port settings:

1. Go to **System > Settings**.
2. Under **Administration Settings**, change the port numbers for HTTP, HTTPS, SSH, and/or Telnet as needed. You can also select **Redirect to HTTPS** in order to avoid HTTP being used for the administrators.
3. Select **Apply**.

When you change the default port number for HTTP, HTTPS, SSH, or Telnet, ensure that the port number is unique. If a conflict exists with a particular port, a warning message will appear.

By default, the GUI disconnects administrative sessions if no activity occurs for five minutes. This prevents someone from using the GUI if the management PC is left unattended.

To change the idle timeout

1. Go to **System > Settings**.
2. In the **Administration Settings** section, enter the time in minutes in the **Idle timeout** field.
3. Select **Apply**.

Password Policy

The FortiGate includes the ability to create a password policy for administrators and IPsec pre-shared keys. With this policy, you can enforce regular changes and specific criteria for a password including:

- minimum length between 8 and 64 characters.
- if the password must contain uppercase (A, B, C) and/or lowercase (a, b, c) characters.
- if the password must contain numbers (1, 2, 3).
- if the password must contain special or non-alphanumeric characters (!, @, #, \$, %, ^, &, *, (, and)).
- where the password applies (admin or IPsec or both).
- the duration of the password before a new one must be specified.

To create a password policy - GUI

1. Go to **System > Settings**.
2. Configure **Password Policy** settings as required.
3. Click **Apply**.

If you add a password policy or change the requirements on an existing policy, the next time that administrator logs into the FortiGate, they are prompted to update their password to meet the new requirements before proceeding to log in.

For information about recovering a lost password and enhancements to the process, see: [Resetting a lost Admin password](#) on the Fortinet Cookbook site.

View Settings

Three settings can change the presentation of information in the GUI: **Language**, **Lines per page**, and **Theme**.

To change the language, go to **System > Settings**. Select the language you want from the **Language** drop-down list: English (the default), French, Spanish, Portuguese, Japanese, Traditional Chinese, Simplified Chinese, or Korean. For best results, you should select the language that is used by the management computer.

To change the number of lines per page displayed in the GUI tables, set **Lines per page** to a value between 20 and 1,000. The default is 50 lines per page.

Five color themes are currently available: Green (the default), Red, Blue, Melongene, and Mariner. To change your theme, select the color from the **Theme** drop-down list.

This is also where you select either **Flow-based** or **Proxy Inspection Mode**. If you select Flow-based mode, then you need to specify if it is **NGFW Profile-based** or **NGFW Policy-based** inspection.

Administrator password retries and lockout time

By default, the FortiGate sets the number of password retries at three, allowing the administrator a maximum of three attempts to log into their account before locking the account for a set amount of time.

Both the number of attempts (`admin-lockout-threshold`) and the wait time before the administrator can try to enter a password again (`admin-lockout-duration`) can be configured within the CLI.

To configure the lockout options:

```
config system global
    set admin-lockout-threshold <failed_attempts>
    set admin-lockout-duration <seconds>
end
```

The default value of `admin-lockout-threshold` is 3 and the range of values is between 1 and 10. The `admin-lockout-duration` is set to 60 seconds by default and the range of values is between 1 and 4294967295 seconds.

Keep in mind that the higher the lockout threshold, the higher the risk that someone may be able to break into the FortiGate.

Example:

To set the `admin-lockout-threshold` to one attempt and the `admin-lockout-duration` to a five minute duration before the administrator can try to log in again, enter the commands:

```
config system global
  set admin-lockout-threshold 1
  set admin-lockout-duration 300
end
```



If the time span between the first failed login attempt and the `admin-lockout-threshold` failed login attempt is less than `admin-lockout-duration`, the lockout will be triggered.

Passwords

Using secure passwords are vital for preventing unauthorized access to your FortiGate. When changing the password, consider the following to ensure better security:

- Do not make passwords that are obvious, such as the company name, administrator names, or other obvious words or phrases.
- Use numbers in place of letters, for example, `passw0rd`.
- Administrator passwords can be up to 64 characters.
- Include a mixture of letters, numbers, and upper and lower case.
- Use multiple words together, or possibly even a sentence, for example `keytothehighway`.
- Use a password generator.
- Change the password regularly and always make the new password unique and not a variation of the existing password, such as changing from `password` to `password1`.
- Make note of the password and store it in a safe place away from the management computer, in case you forget it or ensure that at least two people know the password in the event that one person becomes ill, is away on vacation, or leaves the company. Alternatively, have two different admin logins.

Downgrades will typically maintain the administrator password. If you need to downgrade to FortiOS 4.3, remove the password before the downgrade, then log in after the downgrade and re-configure the password.

Password policy

The FortiGate includes the ability to create a password policy for administrators and IPsec pre-shared keys. With this policy, you can enforce regular changes and specific criteria for a password including:

- minimum length between 8 and 64 characters.
- if the password must contain uppercase (A, B, C) and/or lowercase (a, b, c) characters.
- if the password must contain numbers (1, 2, 3).
- if the password must contain special or non-alphanumeric characters (!, @, #, \$, %, ^, &, *, (, and)).
- where the password applies (admin or IPsec or both).
- the duration of the password before a new one must be specified.

To create a password policy - GUI

1. Go to **System > Settings**.
2. Configure **Password Policy** settings as required.
3. Click **Apply**.

If you add a password policy or change the requirements on an existing policy, the next time that administrator logs into the FortiGate, they are prompted to update their password to meet the new requirements before proceeding to log in.

For information about recovering a lost password and enhancements to the process, see: [Resetting a lost Admin password](#) on the Fortinet Cookbook site.

Firmware

Fortinet periodically updates the FortiGate firmware to include new features and resolve important issues. After you have registered your FortiGate unit, you can download firmware updates from the support web site, <https://support.fortinet.com>.

Before you install any new firmware, be sure to follow the steps below:

- Review the [Release Notes](#) for a new firmware release.
- Review the [Supported Upgrade Paths](#) Sys Admin note on the Fortinet Cookbook site to make sure the upgrade from your current image to the desired new image is supported.
- Backup the current configuration, including local certificates.
- Test the new firmware until you are satisfied that it applies to your configuration.

Installing new firmware without reviewing release notes or testing the firmware may result in changes to settings or unexpected issues.



Only FortiGate admin users and administrators whose access profiles contain system read and write privileges can change the FortiGate firmware.

Backing up the current configuration

You should always back up the configuration before installing new firmware, in case you need to restore your FortiGate configuration.

To create a local backup:

1. Open to the administrator's dropdown menu in the top-right corner of the GUI and select **Configuration > Backup**.
2. Choose either **Local PC** or **USB Disk** to save the configuration file. The USB option will not be available if there is no USB drive in the USB port.
3. If desired, select **Encryption**.
4. Select **OK**.

For more information, see [Configuration Backups](#).

Restoring configuration

Rather than reconfigure the FortiGate manually, it is possible to upload a saved configuration file.

To restore your FortiGate configuration:

1. Open to the administrator's dropdown menu in the top-right corner of the GUI and select **Configuration > Restore**.
2. Choose either **Local PC** or **USB Disk** to restore the configuration file from.
3. Select **Upload** beside **File**.
4. Locate and then select the correct file in the file manager window.
5. If a password was associated with the configuration file, enter it in the **Password** field.
6. Select **OK**.

Troubleshooting

During the installation, some possible errors may occur, but the solutions are usually straightforward.

Error message	Reason and Solution
Configuration file error	<p>This error occurs when attempting to upload a configuration file that is incompatible with the device. This may be due to the configuration file being for a different model or being saved from a different version of firmware.</p> <p>Solution: Upload a configuration file that is for the correct model of FortiGate device and the correct version of the firmware.</p>
Invalid password	<p>When the configuration file is saved, it can be protected by a password. The password entered during the upload process is not matching the one associated with the configuration file.</p> <p>Solution: Use the correct password if the file is password protected.</p>

Downloading firmware

Firmware images for all FortiGate units are available on the Fortinet Customer Support website, <https://support.fortinet.com>.

To download firmware:

1. Log into the site using your user name and password.
2. Go to **Download > Firmware Images**.
3. A list of Release Notes is shown. If you have not already done so, download and review the Release Notes for the firmware you wish to upgrade your FortiGate unit to.
4. Select **Download**.



Firmware can also be downloaded using FTP; however, as FTP is not an encrypted file transferring protocol, HTTPS downloading is recommended.

5. Navigate to the folder for the firmware version you wish to use.
6. Select your FortiGate model from the list. If your unit is a FortiWiFi, the firmware will have a filename starting with 'FWF'.
7. Save the firmware image to your computer.

Testing new firmware

The integrity of firmware images downloaded from Fortinet's support portal can be verified using a file checksum. A file checksum that does not match the expected value indicates a corrupt file. The corruption could be caused by errors in transfer or by file modification. A list of expected checksum values for each build of released code is available on Fortinet's support portal.

Image integrity is also verified when the FortiGate is booting up. This integrity check is done through a cyclic redundancy check (CRC). If the CRC fails, the FortiGate unit will encounter an error during the boot process.

Lastly, firmware images are signed and the signature is attached to the code as it is built. When upgrading an image, the running OS will generate a signature and compare it with the signature attached to the image. If the signatures do not match, the new OS will not load.

Testing before installation

FortiOS lets you test a new firmware image by installing the firmware image from a system reboot and saving it to system memory. After completing this procedure, the FortiGate unit operates using the new firmware image with the current configuration. This new firmware image is not permanently installed. The next time the FortiGate unit restarts, it operates with the originally installed firmware image using the current configuration. If the new firmware image operates successfully, you can install it permanently using the procedure explained in [Testing new firmware on page 117](#).

To use this procedure, you must connect to the CLI using the FortiGate console port and an RJ-45 to DB-9 or null modem cable. This procedure temporarily installs a new firmware image using your current configuration.

For this procedure, you must install a TFTP server that you can connect to from the FortiGate internal interface. The TFTP server should be on the same subnet as the internal interface.

To test the new firmware image:

1. Connect to the CLI using an RJ-45 to DB-9 or null modem cable.
2. Make sure the TFTP server is running.
3. Copy the new firmware image file to the root directory of the TFTP server.
4. Make sure the FortiGate unit can connect to the TFTP server using the `execute ping` command.
5. Enter the following command to restart the FortiGate unit:
`execute reboot`
6. As the FortiGate unit reboots, press any key to interrupt the system startup. As the FortiGate unit starts, a series of system startup messages appears:
Press any key to display configuration menu....
7. Immediately press any key to interrupt the system startup.



You have only three (3) seconds to press any key. If you do not press a key quickly enough, the FortiGate unit reboots and you must log in and repeat the `execute reboot` command.

If you successfully interrupt the startup process, the following messages appears:

```
[G]: Get firmware image from TFTP server.
[F]: Format boot device.
[B]: Boot with backup firmware and set as default
[C]: Configuration and information
[Q]: Quit menu and continue to boot with default firmware.
[H]: Display this list of options.
Enter G, F, Q, or H:
```

8. Type **G** to get the new firmware image from the TFTP server. The following message appears:

```
Enter TFTP server address [192.168.1.168]:
```

9. Type the address of the TFTP server and press **Enter**. The following message appears:

```
Enter Local Address [192.168.1.188]:
```

10. Type an IP address of the FortiGate unit to connect to the TFTP server. The IP address must be on the same network as the TFTP server.



Make sure you do not enter the IP address of another device on this network.

The following message appears:

```
Enter File Name [image.out]:
```

11. Enter the firmware image file name and press **Enter**. The TFTP server uploads the firmware image file to the FortiGate unit and the following appears.

```
Save as Default firmware/Backup firmware/Run image without saving: [D/B/R]
```

12. Type **R**. The FortiGate image is installed to system memory and the FortiGate unit starts running the new firmware image, but with its current configuration.

You can test the new firmware image as required. When done testing, you can reboot the FortiGate unit, and the FortiGate unit will resume using the firmware that was running before you installed the test firmware.

Upgrading the firmware

Installing firmware replaces your current antivirus and attack definitions, along with the definitions included with the firmware release you are installing. After you install new firmware, make sure that antivirus and attack definitions are up to date. You can also use the CLI command `execute update-now` to update the antivirus and attack definitions. For more information, see the [System Administration](#) handbook.



Always remember to back up your configuration before making any changes to the firmware.

Be sure to read the topics on [downloading](#) and [testing](#) firmware before upgrading.

To upgrade the firmware - GUI:

1. Log into the GUI as the admin administrative user.
2. Go to **System > Firmware**.
3. Under **Upload Firmware**, select **Browse** and locate the firmware image file.
4. Select **OK**.

The FortiGate unit uploads the firmware image file, upgrades to the new firmware version, restarts, and displays the FortiGate login. This process takes a few minutes.

To upgrade the firmware - CLI:

Before you begin, ensure you have a TFTP server running and accessible to the FortiGate unit.

1. Make sure the TFTP server is running.
2. Copy the new firmware image file to the root directory of the TFTP server.
3. Log into the CLI.
4. Make sure the FortiGate unit can connect to the TFTP server. You can use the following command to ping the computer running the TFTP server. For example, if the IP address of the TFTP server is 192.168.1.168:

```
execute ping 192.168.1.168
```

5. Enter the following command to copy the firmware image from the TFTP server to the FortiGate unit:

```
execute restore image tftp <filename> <tftp_ipv4>
```

Where `<name_str>` is the name of the firmware image file and `<tftp_ipv4>` is the IP address of the TFTP server. For example, if the firmware image file name is `image.out` and the IP address of the TFTP server is `192.168.1.168`, enter:

```
execute restore image tftp image.out 192.168.1.168
```

The FortiGate unit responds with the message:

```
This operation will replace the current firmware version!  
Do you want to continue? (y/n)
```

6. Type `y`. The FortiGate unit uploads the firmware image file, upgrades to the new firmware version, and restarts. This process takes a few minutes.
7. Reconnect to the CLI.
8. Update antivirus and attack definitions, by entering:

```
execute update-now
```

Reverting to a previous firmware version

The following procedure reverts the FortiGate unit to its factory default configuration and deletes any configuration settings. If you are reverting to a previous FortiOS version, you might not be able to restore the previous configuration from the backup configuration file.



Always remember to back up your configuration before making any changes to the firmware.

To revert to a previous firmware version - GUI:

1. Log into the GUI as the admin user.
2. Go to **System > Firmware**.
3. Under **Upload Firmware**, select **Browse** and locate the firmware image file.
4. Select **OK**.

The FortiGate unit uploads the firmware image file, reverts to the old firmware version, resets the configuration, restarts, and displays the FortiGate login. This process takes a few minutes.

To revert to a previous firmware version - CLI:

Before beginning this procedure, it is recommended that you:

- Backup the FortiGate unit system configuration using the command

```
execute backup config.
```
- Backup the IPS custom signatures using the command

```
execute backup ipsuserdefs sig.
```
- Backup web content and email filtering lists.

To use the following procedure, you must have a TFTP server the FortiGate unit can connect to.

1. Make sure the TFTP server is running.
2. Copy the firmware image file to the root directory of the TFTP server.

3. Log into the FortiGate CLI.
4. Make sure the FortiGate unit can connect to the TFTP server by using the `execute ping` command.
5. Enter the following command to copy the firmware image from the TFTP server to the FortiGate unit:

```
execute restore image tftp <name_str> <tftp_ipv4>
```

Where `<name_str>` is the name of the firmware image file and `<tftp_ipv4>` is the IP address of the TFTP server. For example, if the firmware image file name is `imagev28.out` and the IP address of the TFTP server is `192.168.1.168`, enter:

```
execute restore image tftp image28.out 192.168.1.168
```

The FortiGate unit responds with this message:

```
This operation will replace the current firmware version!
Do you want to continue? (y/n)
```

6. Type `y`. The FortiGate unit uploads the firmware image file. After the file uploads, a message similar to the following appears:

```
Get image from tftp server OK.
Check image OK.
This operation will downgrade the current firmware version!
Do you want to continue? (y/n)
```

7. Type `y`. The FortiGate unit reverts to the old firmware version, resets the configuration to factory defaults, and restarts. This process takes a few minutes.
8. Reconnect to the CLI.
9. To restore your previous configuration, if needed, use the command:

```
execute restore config <name_str> <tftp_ipv4>
```

10. Update antivirus and attack definitions using the command:

```
execute update-now
```

Installing firmware from a system reboot - CLI

In the event that the firmware upgrade does not load properly and the FortiGate unit will not boot, or continuously reboots, it is best to perform a fresh install of the firmware from a reboot using the CLI.

This procedure installs a firmware image and resets the FortiGate unit to default settings. You can use this procedure to upgrade to a new firmware version, revert to an older firmware version, or re-install the current firmware.

To use this procedure, you must connect to the CLI using the FortiGate console port and a RJ-45 to DB-9, or null modem cable. This procedure reverts the FortiGate unit to its factory default configuration.

For this procedure you install a TFTP server that you can connect to from the FortiGate internal interface. The TFTP server should be on the same subnet as the internal interface.

Before beginning this procedure, ensure you backup the FortiGate unit configuration.

If you are reverting to a previous FortiOS version, you might not be able to restore the previous configuration from the backup configuration file.

Installing firmware replaces your current antivirus and attack definitions, along with the definitions included with the firmware release you are installing. After you install new firmware, make sure that antivirus and attack definitions are up to date.

To install firmware from a system reboot:

1. Connect to the CLI using the RJ-45 to DB-9 or null modem cable.
2. Make sure the TFTP server is running.
3. Copy the new firmware image file to the root directory of the TFTP server.
4. Make sure the internal interface is connected to the same network as the TFTP server.
5. To confirm the FortiGate unit can connect to the TFTP server, use the following command to ping the computer running the TFTP server. For example, if the IP address of the TFTP server is 192.168.1.168:

```
execute ping 192.168.1.168
```

6. Enter the following command to restart the FortiGate unit.

```
execute reboot
```

The FortiGate unit responds with the following message:

```
This operation will reboot the system!  
Do you want to continue? (y/n)
```

7. Type **y**. As the FortiGate unit starts, a series of system startup messages appears. When the following messages appears:

```
Press any key to display configuration menu.....
```

Immediately press any key to interrupt the system startup.



You have only three (3) seconds to press any key. If you do not press a key quickly enough, the FortiGate unit reboots and you must log in and repeat the `execute reboot` command.

If you successfully interrupt the startup process, the following messages appears:

```
[G]: Get firmware image from TFTP server.  
[F]: Format boot device.  
[B]: Boot with backup firmware and set as default  
[C]: Configuration and information  
[Q]: Quit menu and continue to boot with default firmware.  
[H]: Display this list of options.  
Enter G, F, Q, or H:
```

8. Type **G** to get to the new firmware image from the TFTP server. The following message appears:

```
Enter TFTP server address [192.168.1.168]:
```

9. Type the address of the TFTP server and press **Enter**. The following message appears:

```
Enter Local Address [192.168.1.188]:
```

10. Type an IP address the FortiGate unit can use to connect to the TFTP server. The IP address can be any IP address that is valid for the network to which the interface is connected.



Make sure you do not enter the IP address of another device on this network.

The following message appears:

```
Enter File Name [image.out]:
```

11. Enter the firmware image filename and press **Enter**.

The TFTP server uploads the firmware image file to the FortiGate unit and a message similar to the following

appears:

```
Save as Default firmware/Backup firmware/Run image without saving: [D/B/R]
```

12. Type **D**. The FortiGate unit installs the new firmware image and restarts. The installation might take a few minutes to complete.

Restore from a USB key - CLI

1. Log into the CLI.
2. Enter the following command to restore an unencrypted configuration file:

```
exec restore image usb <filename>
```

If your configuration file was encrypted, enter the following command:

```
execute restore config usb-mode <password>
```

The FortiGate unit responds with the following message:

```
This operation will replace the current firmware version!  
Do you want to continue? (y/n)
```

3. Type **y**.

Controlled upgrade

Using a controlled upgrade, you can upload a new version of the FortiOS firmware to a separate partition in the FortiGate memory for later upgrade. The FortiGate unit can also be configured so that when it is rebooted, it will automatically load the new firmware (CLI only). Using this option, you can stage a number of FortiGate units to do an upgrade simultaneously to all devices using FortiManager or script.

To load the firmware for later installation - CLI:

```
execute restore secondary-image {ftp | tftp | usb} <filename_str>
```

To set the FortiGate unit so that when it reboots, the new firmware is loaded, use the CLI command . . .

```
execute set-next-reboot {primary | secondary}
```

where {primary | secondary} is the partition with the preloaded firmware.

Configuration backups

Once you successfully configure the FortiGate, it is extremely important that you backup the configuration. In some cases, you may need to reset the FortiGate to factory defaults or perform a TFTP upload of the firmware, which will erase the existing configuration. In these instances, the configuration on the device will have to be recreated, unless a backup can be used to restore it. You should also backup the local certificates, as the unique SSL inspection CA and server certificates that are generated by your FortiGate by default are not saved in a system backup.

It is also recommended that you backup the configuration after *any* future changes are made, to ensure you have the most current configuration available. Also, backup the configuration before any upgrades of the FortiGate's firmware. Should anything happen to the configuration during the upgrade, you can easily restore the saved configuration.

Always backup the configuration and store it on the management computer or off-site. You have the option to save the configuration file to various locations including the local PC, USB key, FTP, and TFTP server. The last two are configurable through the CLI only.

If you have VDOMs, you can back up the configuration of the entire FortiGate or only a specific VDOM. Note that if you are using FortiManager or FortiCloud, full backups are performed and the option to backup individual VDOMs will not appear.

Backing up the configuration using the GUI

1. Click on **admin** in the upper right-hand corner of the screen and select **Configuration > Backup**.
2. Direct the backup to your **Local PC** or to a **USB Disk**.
The **USB Disk** option will be grayed out if no USB drive is inserted in the USB port. You can also backup to the FortiManager using the CLI.
3. If VDOMs are enabled, indicate whether the scope of the backup is for the entire FortiGate configuration (**Global**) or only a specific VDOM configuration (**VDOM**).
4. If backing up a VDOM configuration, select the VDOM name from the list.
5. Select **Encryption**.
Encryption must be enabled on the backup file to back up VPN certificates.
6. Enter a password and enter it again to confirm it. You will need this password to restore the file.
7. Select **OK**.
8. The web browser will prompt you for a location to save the configuration file. The configuration file will have a .conf extension.

Backing up the configuration using the CLI

Use one of the following commands:

```
execute backup config management-station <comment>
```

or:

```
execute backup config usb <backup_filename> [<backup_password>]
```

or for FTP, note that port number, username are optional depending on the FTP site:

```
execute backup config ftp <backup_filename> <ftp_server> [<port>] [<user_name>]
[<password>]
```

or for TFTP:

```
execute backup config tftp <backup_filename> <tftp_servers> <password>
```

Use the same commands to backup a VDOM configuration by first entering the commands:

```
config vdom
edit <vdom_name>
```

Backup and restore the local certificates

This procedure exports a server (local) certificate and private key together as a password protected PKCS12 file. The export file is created through a customer-supplied TFTP server. Ensure that your TFTP server is running and accessible to the FortiGate before you enter the command.

To back up the local certificates:

Connect to the CLI and use the following command:

```
execute vpn certificate local export tftp <cert_name> <filename> <tftp_ip>
```

where:

- <cert_name> is the name of the server certificate.
- <filename> is a name for the output file.
- <tftp_ip> is the IP address assigned to the TFTP server host interface.

To restore the local certificates - GUI:

1. Move the output file from the TFTP server location to the management computer.
2. Go to **System > Certificates** and select **Import**.
3. Select the appropriate type of certificate from the dropdown menu and fill in any required fields.
4. Select **Upload**. Browse to the location on the management computer where the exported file has been saved, select the file and select **Open**.
5. If required, enter the **Password** needed to upload the exported file.
6. Select **OK**.

To restore the local certificates - CLI:

Connect to the CLI and use the following command:

```
execute vpn certificate local import tftp <filename> <tftp_ip>
```

Backup and restore a configuration file using SCP

You can use secure copy protocol (SCP) to download the configuration file from the FortiGate as an alternative method of backing up the configuration file or an individual VDOM configuration file. This is done by enabling SCP for an administrator account and enabling SSH on a port used by the SCP client application to connect to the FortiGate. SCP is enabled using the CLI commands:

```
config system global
set admin-scp enable
end
```

Use the same commands to backup a VDOM configuration by first entering the commands:

```
config global
    set admin-scp enable
end
config vdom
    edit <vdom_name>
```

Enable SSH access on the interface

SCP uses the SSH protocol to provide secure file transfer. The interface you use for administration must allow SSH access.

To enable SSH - GUI:

1. Go to **Network > Interfaces**.
2. Select the interface you use for administrative access and select **Edit**.
3. In the **Administrative Access** section, select **SSH**.
4. Select **OK**.

To enable SSH - CLI:

```
config system interface
    edit <interface_name>
        set allowaccess ping https ssh
    end
```



When adding to, or removing a protocol, you must type the entire list again. For example, if you have an access list of HTTPS and SSH, and you want to add PING, typing:

```
set allowaccess ping
```

will only set PING. To add all three, you must type:

```
set allowaccess https ssh ping
```

Using the SCP client

The FortiGate downloads the configuration file as `sys_conf`. Use the following syntax to download the file:

Linux

```
scp admin@<FortiGate_IP>:fgt-config <location>
```

Windows

```
pscp admin@<FortiGate_IP>:fgt-config <location>
```

The following examples show how to download the configuration file from a FortiGate-100D, at IP address 172.20.120.171, using Linux and Windows SCP clients.

Linux client example

To download the configuration file to a local directory called `~/config`, enter the following command:

```
scp admin@172.20.120.171:fgt-config ~/config
```

Enter the admin password when prompted.

Windows client example

To download the configuration file to a local directory called c:\config, enter the following command in a Command Prompt window:

```
pscp admin@172.20.120.171:fgt-config c:\config
```

Enter the admin password when prompted.

SCP public-private key authentication

SCP authenticates itself to the FortiGate in the same way as an administrator using SSH accesses the CLI. Instead of using a password, you can configure the SCP client and the FortiGate with a public-private key pair.

To configure public-private key authentication:

1. Create a public-private key pair using a key generator compatible with your SCP client.
2. Save the private key to the location on your computer where your SSH keys are stored.
This step depends on your SCP client. The Secure Shell key generator automatically stores the private key.

3. Copy the public key to the FortiGate using the CLI commands:

```
config system admin
  edit admin
    set ssh-public-key1 "<key-type> <key-value>"
  end
```

<key-type> must be the ssh-dss for a DSA key or ssh-rsa for an RSA key. For the <key-value>, copy the public key data and paste it into the CLI command.

If you are copying the key data from Windows Notepad, copy one line at a time and ensure that you paste each line of key data at the end of the previously pasted data. Also:

- Do not copy the end-of-line characters that appear as small rectangles in Notepad.
- Do not copy the ---- BEGIN SSH2 PUBLIC KEY ---- or Comment: "[2048-bit dsa,...]" lines.
- Do not copy the ---- END SSH2 PUBLIC KEY ---- line.

4. Type the closing quotation mark and press **Enter**.

Your SCP client can now authenticate to the FortiGate based on SSH keys rather than the administrator password.

Restoring a configuration using SCP

To restore the configuration using SCP, use the commands:

```
scp <local_file> <admin_user>@<FGT_IP>:fgt-restore-config
```

To use this command/method of restoring the FortiGate configuration, you need to log in as the "admin" administrator.

Restoring a configuration

Should you need to restore a configuration file, use the following steps:

To restore the FortiGate configuration - GUI:

1. Click on **admin** in the upper right-hand corner of the screen and select **Configuration > Restore**.
2. Identify the source of the configuration file to be restored : your **Local PC** or a **USB Disk**.
The **USB Disk** option will be grayed out if no USB drive is inserted in the USB port. You can restore from the FortiManager using the CLI.
3. Enter the path and file name of the configuration file, or select **Browse** to locate the file.
4. Enter a password if required.
5. Select **Restore**.

To back up the FortiGate configuration - CLI:

```
execute restore config management-station normal 0
```

or:

```
execute restore config usb <filename> [<password>]
```

or for FTP, note that port number, username are optional depending on the FTP site:

```
execute backup config ftp <backup_filename> <ftp_server> [<port>] [<user_name>] [<password>]
```

or for TFTP:

```
execute backup config tftp <backup_filename> <tftp_server> <password>
```

The FortiGate will load the configuration file and restart. Once the restart has completed, verify that the configuration has been restored.

Configuration revision

You can manage multiple versions of configuration files on models that have a 512MB flash memory and higher. Revision control requires either a configured central management server or the local hard drive, if your FortiGate has this feature. Typically, configuration backup to local drive is not available on lower-end models.

The central management server can either be a FortiManager unit or FortiCloud.

If central management is not configured on your FortiGate unit, a message appears instructing you to either:

- Enable central management, **or**
- obtain a valid license.

When revision control is enabled on your FortiGate unit, and configuration backups have been made, a list of saved revisions of those backed-up configurations appears.

Configuration revisions are viewed by clicking on **admin** in the upper right-hand corner of the screen and selecting **Configuration > Revisions**.

Restore factory defaults

There may be a need to reset the FortiGate to its original defaults; for example, to begin with a fresh configuration. There are two options when restoring factory defaults. The first resets the entire device to the original out-of-the-box configuration.

You can reset using the CLI by entering the command:

```
execute factoryreset
```

When prompted, type `y` to confirm the reset.

Alternatively, in the CLI you can reset the factory defaults but retain the interface and VDOM configuration. Use the following command:

```
execute factoryreset2
```


FortiGuard

The FortiGuard Distribution Network (FDN) of servers provides updates to antivirus, antispyware, and IPS definitions to your FortiGate. FortiGuard Subscription Services provides comprehensive Unified Threat Management (UTM) security solutions to enable protection against content and network level threats.

The FortiGuard team can be found around the globe, monitoring virus, spyware and vulnerability activities. As vulnerabilities are found, signatures are created and pushed to the subscribed FortiGates. The Global Threat Research Team enables Fortinet to deliver a combination of multi-layered security intelligence and provide true zero-day protection from new and emerging threats. The FortiGuard Network has data centers around the world located in secure, high availability locations that automatically deliver updates to the Fortinet security platforms to protect the network with the latest information.

FortiGuard provides a number of services to monitor world-wide activity and provide the best possible security, including:

- **Intrusion Prevention System (IPS)** - IPS uses a customizable database of more than 4000 known threats to stop attacks that evade conventional firewall defenses. It also provides behavior-based heuristics, enabling the system to recognize threats when no signature has yet been developed. It also provides more than 1000 application identity signatures for complete application control.
- **Application Control** - Application Control allows you to identify and control applications on networks and endpoints regardless of port, protocol, and IP address used. It gives you unmatched visibility and control over application traffic, even traffic from unknown applications and sources. Application Control is a free FortiGuard service and the database for Application Control signatures is separate from the IPS database (Botnet Application signatures are still part of the IPS signature database since these are more closely related with security issues and less about application detection). Application Control signature database information is displayed under the **System > FortiGuard** page in the FortiCare section.



Please note that while the Application Control profile can be used for free, signature database updates require a valid FortiGuard subscription.

- **AntiVirus** - The FortiGuard AntiVirus Service provides fully automated updates to ensure protection against the latest content level threats. It employs advanced virus, spyware, and heuristic detection engines to prevent both new and evolving threats from gaining access to your network and protects against vulnerabilities.
- **Web Filtering** - Web Filtering provides Web URL filtering to block access to harmful, inappropriate, and dangerous web sites that may contain phishing/pharming attacks, malware such as spyware, or objectionable content that can expose your organization to legal liability. Based on automatic research tools and targeted research analysis, real-time updates enable you to apply highly-granular policies that filter web access based on six major categories and nearly 80 micro-categories, over 45 million rated web sites, and more than two billion web pages - all continuously updated.
- **Email Filtering** - The FortiGuard Antispam Service uses both a sender IP reputation database and a spam signature database, along with sophisticated spam filtering tools on Fortinet appliances and agents, to detect and block a wide range of spam messages. Updates to the IP reputation and spam signature databases are provided continuously via the FDN.
- **Messaging Services** - Messaging Services allow a secure email server to be automatically enabled on your FortiGate to send alert email or send email authentication tokens. With the SMS gateway, you can enter phone numbers where the FortiGate will send the SMS messages. Note that depending on your carrier, there may be a slight time delay on receiving messages.

- **DNS and DDNS** - The FortiGuard DNS and DDNS services provide an efficient method of DNS lookups once subscribed to the FortiGuard network. This is the default option. The FortiGate connects automatically to the FortiGuard DNS server. If you do not register, you need to configure an alternate DNS server.

Configure the DDNS server settings using the CLI command:

```
config system fortiguard
    set ddns-server-ip
    set ddns-server-port
end
```

Support contract and FortiGuard subscription services

The FDN support **Contract** is available under **System > FortiGuard**.

The License Information area displays the status of your FortiGate's support contract.

You can also manually update the AntiVirus and IPS engines.

Verifying your connection to FortiGuard

If you are not getting FortiGuard web filtering or antispam services, there are a few things to verify that communication to the FDN is working. Before any troubleshooting, ensure that the FortiGate has been registered and subscribed to the FortiGuard services.

Verification - GUI:

The simplest method to check that the FortiGate is communicating with the FDN, is to check the **License Information** dashboard widget. Any subscribed services should have a green check mark beside them indicating that connections are successful. Any other icon indicates a problem with the connection, or you are not subscribed to the FortiGuard services.

You can also view the FortiGuard connection status by going to **System > FortiGuard**.

Verification - CLI:

You can also use the CLI to see what FortiGuard servers are available to your FortiGate. Use the following CLI command to ping the FDN for a connection:

```
execute ping guard.fortinet.net
```

You can also use the following diagnose command to find out what FortiGuard servers are available:

```
diagnose debug rating
```

From this command, you will see output similar to the following:

```
Locale : english
License : Contract
Expiration : Sun Jul 24 20:00:00 2011
Hostname : service.fortiguard.net
== Server List (Tue Nov 2 11:12:28 2010) ==
```

IP	Weight	RTT	Flags	TZ	Packets	Curr	Lost	Total	Lost
69.20.236.180	0	10		-5	77200	0	42		
69.20.236.179	0	12		-5	52514	0		34	
66.117.56.42	0	32		-5	34390	0		62	
80.85.69.38	50	164		0	34430	0		11763	
208.91.112.194	81	223	D	-8	42530	0		8129	
216.156.209.26	286	241	DI	-8	55602	0		21555	

An extensive list of servers are available. Should you see a list of three to five available servers, the FortiGuard servers are responding to DNS replies to service FortiGuard.net, but the INIT requests are not reaching FDS services on the servers.

The rating flags indicate the server status:

D	Indicates the server was found via the DNS lookup of the hostname. If the hostname returns more than one IP address, all of them will be flagged with 'D' and will be used first for INIT requests before falling back to the other servers.
I	Indicates the server to which the last INIT request was sent.
F	The server has not responded to requests and is considered to have failed.
T	The server is currently being timed.

The server list is sorted first by weight and then the server with the smallest RTT is put at the top of the list, regardless of weight. When a packet is lost, it will be resent to the next server in the list.

The weight for each server increases with failed packets and decreases with successful packets. To lower the possibility of using a distant server, the weight is not allowed to dip below a base weight, which is calculated as the difference in hours between the FortiGate and the server, multiplied by 10. The further away the server, the higher its base weight and the lower in the list it will appear.

Port assignment

The FortiGate contacts FDN for the latest list of FDN servers by sending UDP packets with typical source ports of 1027 or 1031, and destination port 8888. The FDN reply packets have a destination port of 1027 or 1031.

If your ISP blocks UDP packets in this port range, the FortiGate cannot receive the FDN reply packets. As a result, the FortiGate will not receive the complete FDN server list.

If your ISP blocks the lower range of UDP ports (around 1024), you can configure your FortiGate to use higher-numbered ports, using the CLI command:

```
config system global
    set ip-src-port-range <start port>-<end port>
end
```

where the <start port> and <end port> are numbers ranging of 1024 to 25000.

For example, you could configure the FortiGate to not use ports lower than 2048 or ports higher than the following range:

```
config system global
    set ip-src-port-range 2048-20000
end
```

Trial and error may be required to select the best source port range. You can also contact your ISP to determine the best range to use. Push updates might be unavailable if:

- there is a NAT device installed between the unit and the FDN, and/or
- your unit connects to the Internet using a proxy server.

Configuring AntiVirus and IPS options

Go to **System > FortiGuard**, and scroll down to the **AntiVirus & IPS Updates** section to configure the antivirus and IPS options for connecting and downloading definition files.

Accept push updates	Select to allow updates to be sent automatically to your FortiGate. New definitions will be added as soon as they are released by FortiGuard.
Use override push	<p>Appears only if Accept push updates is enabled.</p> <p>Enable to configure an override server if you cannot connect to the FDN or if your organization provides updates using their own FortiGuard server. Once enabled, enter the following:</p> <ul style="list-style-type: none"> • Enter the IP address and port of the NAT device in front of your FortiGate. FDS will connect to this device when attempting to reach the FortiGate. • The NAT device must be configured to forward the FDS traffic to the FortiGate on UDP port 9443.
Scheduled Updates	<p>Enable for updates to be sent to your FortiGate at a specific time. For example, to minimize traffic lag times, you can schedule the update to occur on weekends or after work hours.</p> <p>Note that a schedule of once a week means any urgent updates will not be pushed until the scheduled time. However, if there is an urgent update required, select the Update Now button.</p>
Improve IPS quality	Enable to help Fortinet maintain and improve IPS signatures. The information sent to the FortiGuard servers when an attack occurs can be used to keep the database current as variants of attacks evolve.
Use extended IPS signature package	Regular IPS database protects against the latest common and in-the-wild attacks. Extended IPS database includes protection from legacy attacks.
Update AV & IPS Definitions	Select to manually initiate an FDN update.

Manual updates

To manually update the signature definitions file, you need to first go to the Support web site at <https://support.fortinet.com>. Once logged in, select **Download > FortiGuard Service Updates**. The browser will present you the most current IPS and AntiVirus signature definitions which you can download.

Once downloaded to your computer, log into the FortiGate to load the definition file.

To load the definition file onto the FortiGate:

1. Go to **System > FortiGuard**.
2. In the **License Information** table, select the **Upgrade Database** link in either the **Application Control Signature**, **IPS**, or **AntiVirus** row.
3. In the pop-up window, select **Upload** and locate the downloaded file and select **Open**.

The upload may take a few minutes to complete.

Automatic updates

The FortiGate can be configured to request updates from FDN on a scheduled basis, or via push notification.

Scheduling updates

Scheduling updates ensures that the virus and IPS definitions are downloaded to your FortiGate on a regular basis, ensuring that you do not forget to check for the definition files yourself.

Updating definitions can cause a very short disruption in traffic currently being scanned while the FortiGate unit applies the new signature database. Ideally, schedule updates during off-peak hours, such as evenings or weekends, when network usage is minimal, to ensure that the network activity will not suffer from the added traffic of downloading the definition files.

To enable scheduled updates - GUI:

1. Go to **System > FortiGuard** and scroll down to **AntiVirus & IPS Updates**.
2. Enable **Scheduled Updates**.
3. Select the frequency of updates.
4. Select **Apply**.

To enable scheduled updates - CLI:

```
config system autoupdate schedule
  set status enable
  set frequency {every | daily | weekly}
  set time <hh:mm>
  set day <day_of_week>
end
```

Push updates

Push updates enable you to get immediate updates when new viruses or intrusions have been discovered and new signatures created. This ensures that the latest signature will be sent to the FortiGate as soon as possible.

When a push notification occurs, the FortiGuard server sends a notice to the FortiGate that there is a new signature definition file available. The FortiGate then initiates a download of the definition file, similar to the scheduled update.

To ensure maximum security for your network, you should have a scheduled update as well as enable the push update, in case an urgent signature is created, and your cycle of the updates only occurs weekly.

To enable push updates - GUI:

1. Go to **System > FortiGuard** and scroll down to **AntiVirus & IPS Updates**.
2. Enable **Accept push updates**.
3. Select **Apply**.

To enable push updates - CLI:

```
config system autoupdate push-update
  set status enable
```

end

Push IP override

If the FortiGate is behind another NAT device (or another FortiGate), to ensure it receives the push update notifications, you need to use an override IP address for the notifications. To do this, you create a virtual IP to map to the external port of the NAT device.

Generally speaking, if there are two FortiGate devices, the following steps need to be completed on the FortiGate NAT device to ensure the FortiGate on the internal network receives the updates:

- Add a port forwarding virtual IP to the FortiGate NAT device that connects to the Internet by going to **Policy & Objects > Virtual IPs**.
- Add a security policy to the FortiGate NAT device that connects to the Internet that includes the port forwarding virtual IP.
- Configure the FortiGate on the internal network with an override push IP and port.

On the FortiGate internal device, the virtual IP is entered as the **Use push override** IP address.

To enable push update override- GUI:

1. Go to **System > FortiGuard** and scroll down to **AntiVirus & IPS Updates**.
2. Enable **Accept push updates**.
3. Enable **Use override push**.
4. Enter the virtual IP address configured on the NAT device.
5. Select **Apply**.

To enable push updates - CLI:

```
config system autoupdate push-update
    set status enable
    set override enable
    set address <vip_address>
end
```

Sending malware statistics to FortiGuard

To support following malware trends and making zero-day discoveries, FortiGate units send encrypted statistics to FortiGuard about IPS, Application Control, and AntiVirus events detected by the FortiGuard services running on your FortiGate. FortiGuard uses the statistics collected to achieve a balance between performance and security effectiveness by moving inactive signatures to an extended signature database.

The statistics include some non-personal information that identifies your FortiGate and its country. The information is never shared with external parties. You can choose to disable the sharing of this information by entering the following CLI command:

```
config system global
    set fds-statistics disable
end
```

Configuring web filtering and email filtering options

Go to **System > FortiGuard**, and scroll down to **Filtering** to set the size of the caches and ports.

Web Filter Cache	Set the Time To Live (TTL) value. This is the number of seconds the FortiGate will store a blocked IP or URL locally, saving time and network access traffic, checking the FortiGuard server. Once the TTL has expired, the FortiGate will contact an FDN server to verify a web address. The TTL must be between 300 and 86400 seconds.
Anti-Spam Cache	Set the TTL value (see above).
FortiGuard Filtering Port	Select the port assignments for contacting the FortiGuard servers.
Filtering Service Availability	Indicates the status of the filtering service. Select Check Again if the filtering service is not available.
Request re-evaluation of a URL's category	Select to re-evaluate a URL's category rating on the FortiGuard Web Filter service.

Email filtering

The FortiGuard data centers monitor and update email databases of known spam sources. With FortiGuard Anti-Spam filtering enabled, the FortiGate verifies incoming email sender addresses and IPs against the database, and takes the necessary actions as defined within the antivirus profiles.

Spam source IP addresses can also be cached locally on the FortiGate, providing a quicker response time, while easing load on the FortiGuard servers, aiding in a quicker response time for less common email address requests.

By default, the anti-spam cache is enabled. The cache includes a TTL value, which is the amount of time an email address will stay in the cache before expiring. You can change this value to shorten or extend the time between 5 and 1,440 minutes.

To modify the antispam cache TTL - GUI:

1. Go to **System > FortiGuard**.
2. Under **Filtering**, enable **Anti-Spam Cache**.
3. Enter the TTL value in minutes.
4. Select **Apply**.

To modify the Anti-Spam filter TTL - CLI:

```
config system fortiguard
    set antispam-cache-ttl <integer>
end
```

Further antispam filtering options can be configured to block, allow, or quarantine specific email addresses. These configurations are available through the **Security Profiles > Anti-Spam** menu. For more information, see the [Security Profiles](#) handbook chapter.

Online security tools

The FortiGuard online center provides a number of online security tools, including but not limited to:

- **URL lookup** — By entering a website address, you can see if it has been rated and what category and classification it is filed as. If you find your website or a site you commonly go to has been wrongly categorized, you can use this page to request that the site be re-evaluated.
<https://fortiguard.com/webfilter>
- **Threat Encyclopedia** — Browse the Fortiguard Labs extensive encyclopedia of threats. Search for viruses, botnet C&C, IPS, endpoint vulnerabilities, and mobile malware.
<https://www.fortiguard.com/encyclopedia>
- **Application Control** — Browse the Fortiguard Labs extensive encyclopedia of applications.
<https://fortiguard.com/appcontrol>

FortiCloud

FortiCloud is a hosted security management and log retention service for FortiGate devices. It gives you centralized reporting, traffic analysis, configuration management, and log retention without the need for additional hardware or software.

FortiCloud offers a wide range of features:

- **Simplified central management** — FortiCloud provides a central web-based management console to manage individual or aggregated FortiGate and FortiWiFi devices. Adding a device to the FortiCloud management subscription is straightforward. FortiCloud has detailed traffic and application visibility across the whole network.
- **Hosted log retention with large default storage allocated** — Log retention is an integral part of any security and compliance program but administering a separate storage system is burdensome. FortiCloud takes care of this automatically and stores the valuable log information in the cloud. Each device is allowed up to 200GB of log retention storage. Different types of logs can be stored including Traffic, System Events, Web, Applications, and Security Events.
- **Monitoring and alerting in real time** — Network availability is critical to a good end-user experience. FortiCloud enables you to monitor your FortiGate network in real time with different alerting mechanisms to pinpoint potential issues. Alerting mechanisms can be delivered via email.
- **Customized or pre-configured reporting and analysis tools** — Reporting and analysis are your eyes and ears into your network's health and security. Pre-configured reports are available, as well as custom reports that can be tailored to your specific reporting and compliance requirements. For example, you may want to look closely at application usage or website violations. The reports can be emailed as PDFs and can cover different time periods.
- **Maintain important configuration information uniformly** — The correct configuration of the devices within your network is essential to maintaining an optimum performance and security posture. In addition, maintaining the correct firmware (operating system) level allows you to take advantage of the latest features.
- **Service security** — All communication (including log information) between the devices and the clouds is encrypted. Redundant data centers are always used to give the service high availability. Operational security measures have been put in place to make sure your data is secure — only you can view or retrieve it.

Registration and activation



Before you can activate a FortiCloud account, you must first register your device.

FortiCloud accounts can be registered manually through the FortiCloud website, <https://www.forticloud.com>, but you can easily register and activate your account directly from your FortiGate.

Activating your FortiCloud account

1. On your device's dashboard, in the **FortiCloud** widget, select the **Activate** button in the status field.
2. A dialogue asking you to register your FortiCloud account appears. Select **Create Account**, enter your information, view and accept the terms and conditions, and select **OK**.
3. A second dialogue window appears, asking you to enter your information to confirm your account. This sends a confirmation email to your registered email. The dashboard widget then updates to show that confirmation is required.
4. Open your email, and follow the confirmation link it contains.

Results

A FortiCloud page will open, stating that your account has been confirmed. The Activation Pending message on the dashboard will change to state the type of account you have ('1GB Free' or '200GB Subscription'), and will provide a link to the FortiCloud portal.

Enabling logging to FortiCloud

1. Go to **Log & Report > Log Settings**.
2. Enable **Send Logs to FortiCloud**.
3. Select **Test Connectivity** to ensure that your FortiGate can connect to the registered FortiCloud account.
4. Scroll down to **GUI Preferences**, set **Display Logs/FortiView From**, to see FortiCloud logs within the FortiGate's GUI.

Logging into the FortiCloud portal

Once logging has been configured and you have registered your account, you can log into the FortiCloud portal and begin viewing your logging results. There are two methods to reach the FortiCloud portal:

- If you have direct networked access to the FortiGate, you can simply open your **Dashboard** and check the **License Information** widget. Next to the current FortiCloud connection status will be a link to reach the FortiCloud Portal.
- If you do not currently have access to the FortiGate's interface, you can visit the FortiCloud website (<https://forticloud.com>) and log in remotely, using your email and password. It will ask you to confirm the FortiCloud account you are connecting to and then you will be granted access. Connected devices can be remotely configured using the Scripts page in the Management Tab, useful if an administrator may be away from the unit for a long period of time.

Cloud sandboxing

FortiCloud can be used for automated sample tracking, or sandboxing, for files from a FortiGate. This allows suspicious files to be sent to be inspected without risking network security. If the file exhibits risky behavior, or is found to contain a virus, a new virus signature is created and added to the FortiGuard antivirus signature database.

Cloud sandboxing is configured by going to **Security Fabric > Settings**. After enabling **Sandbox Inspection**, select the **FortiSandbox type**.

Sandboxing results are shown in a new tab called **AV Submissions** in the FortiCloud portal. This tab only appears after a file has been sent for sandboxing.

For more information about FortiCloud, see the [FortiCloud documentation](#).

Troubleshooting your FortiGate installation

If your FortiGate does not function as desired after installation, try the following troubleshooting tips:

1. Check for equipment issues

Verify that all network equipment is powered on and operating as expected. Refer to the QuickStart Guide for information about connecting your FortiGate to the network. You will also find detailed information about the FortiGate LED indicators.

The FortiGate has multiple LED lights on the faceplate. Check the [FortiGate LED specifications](#) guide to verify if the lights indicate any problems with your FortiGate itself or with connections between your FortiGate and other devices.

2. Check the physical network connections

Check the cables used for all physical connections to ensure that they are fully connected and do not appear damaged, and make sure that each cable connects to the correct device and the correct Ethernet port on that device.

3. Verify that you can connect to the internal IP address of the FortiGate

Connect to the GUI from the FortiGate's internal interface by browsing to its IP address. From the PC, try to ping the internal interface IP address; for example, `ping 192.168.1.99`.

If you cannot connect to the internal interface, verify the IP configuration of the PC. If you can ping the interface but can't connect to the GUI, check the settings for administrative access on that interface. Alternatively, use SSH to connect to the CLI, and then confirm that HTTPS has been enabled for Administrative Access on the interface.

4. Check the FortiGate interface configurations

Check the configuration of the FortiGate interface connected to the internal network (under **Network > Interfaces**) and check that **Addressing mode** is set to the correct mode.

5. Verify the security policy configuration

Go to **Policy & Objects > IPv4 Policy** and verify that the internal interface to Internet-facing interface security policy has been added and is located near the top of the policy list. Check the **Active Sessions** column to ensure that traffic has been processed (if this column does not appear, right-click on the table header and select **Active Sessions**).

If you are using NAT/Route mode, check the configuration of the policy to make sure that **NAT** is enabled and that **Use Outgoing Interface Address** is selected.

6. Verify the static routing configuration

Go to **Network > Static Routes** and verify that the default route is correct. Go to **Monitor > Routing Monitor** and verify that the default route appears in the list as a static route. Along with the default route, you should see two routes shown as **Connected**, one for each connected FortiGate interface.

7. Verify that you can connect to the Internet-facing interface's IP address

Ping the IP address of the Internet-facing interface of your FortiGate. If you cannot connect to the interface, the FortiGate is not allowing sessions from the internal interface to Internet-facing interface. Verify that **PING** has been enabled for **Administrative Access** on the interface.

8. Verify that you can connect to the gateway provided by your ISP

Ping the default gateway IP address from a PC on the internal network. If you cannot reach the gateway, contact your ISP to verify that you are using the correct gateway.

9. Verify that you can communicate from the FortiGate to the Internet

Access the FortiGate CLI and use the command `execute ping 8.8.8.8`. You can also use the `execute traceroute 8.8.8.8` command to troubleshoot connectivity to the Internet.

10. Verify the DNS configurations of the FortiGate and the PCs

Check for DNS errors by pinging or using traceroute to connect to a domain name; for example: `ping`

`www.fortinet.com`.

If the name cannot be resolved, the FortiGate or PC cannot connect to a DNS server and you should confirm that the DNS server IP addresses are present and correct.

11. Confirm that the FortiGate can connect to the FortiGuard network

Once the FortiGate is on your network, you should confirm that it can reach the FortiGuard network.

First, check the **License Information** widget to make sure that the status of all FortiGuard services matches the services that you have purchased.

Go to **System > FortiGuard**. Scroll down to **Filtering Services Availability** and select **Check Again**. After a minute, the GUI should indicate a successful connection.

Verify that your FortiGate can resolve and reach FortiGuard at *service.fortiguard.net* by pinging the domain name. If you can reach this service, you can then verify the connection to FortiGuard servers by running the command `diagnose debug rating`. This displays a list of FortiGuard IP gateways you can connect to, as well as the following information:

- **Weight:** Based on the difference in time zone between the FortiGate and this server
- **RTT:** Return trip time
- **Flags:** D (IP returned from DNS), I (Contract server contacted), T (being timed), F (failed)
- **TZ:** Server time zone
- **Curr Lost:** Current number of consecutive lost packets
- **Total Lost:** Total number of lost packets

12. Consider changing the MAC address of your external interface

Some ISPs do not want the MAC address of the device connecting to their network cable to change. If you have added a FortiGate to your network, you may have to change the MAC address of the Internet-facing interface using the following CLI command:

```
config system interface
  edit <interface>
    set macaddr <xx:xx:xx:xx:xx:xx>
  end
end
```

13. Check the FortiGate bridge table (transparent mode)

When a FortiGate is in transparent mode, the unit acts like a bridge sending all incoming traffic out on the other interfaces. The bridge is between interfaces on the FortiGate unit. Each bridge listed is a link between interfaces. Where traffic is flowing between interfaces, you expect to find bridges listed. If you are having connectivity issues and there are no bridges listed, that is a likely cause. Check for the MAC address of the interface or device in question.

To list the existing bridge instances on the FortiGate, use the following CLI command:

```
diagnose netlink brctl name host root.b
show bridge control interface root.b host.
fdb: size=2048, used=25, num=25, depth=1
Bridge root.b host table
port no device devname mac addr ttl attributes
3 4 wan1 00:09:0f:cb:c2:77 88
3 4 wan1 00:26:2d:24:b7:d3 0
3 4 wan1 00:13:72:38:72:21 98
4 3 internal 00:1a:a0:2f:bc:c6 6
1 6 dmz 00:09:0f:dc:90:69 0 Local Static
3 4 wan1 c4:2c:03:0d:3a:38 81
3 4 wan1 00:09:0f:15:05:46 89
3 4 wan1 c4:2c:03:1d:1b:10 0
2 5 wan2 00:09:0f:dc:90:68 0 Local Static
```

14. Use FortiExplorer if you can't connect to the FortiGate over Ethernet

If you can't connect to the FortiGate GUI or CLI, you may be able to connect using FortiExplorer. Refer to the

QuickStart Guide or see the section on [FortiExplorer](#) for more details.

15. Either reset the FortiGate to factory defaults or contact Fortinet Support for assistance

To reset the FortiGate to factory defaults, use the CLI command `execute factoryreset`. When prompted, type `y` to confirm the reset.

You can also contact Fortinet Support for assistance. Read the following article found on the Fortinet Cookbook website: [How to work with Fortinet Support](#) to understand what type of support is available and to determine which level of support is right for you. For further assistance, visit the [Fortinet Support](#) website.

Resources

Here's a list of some resources you can check out next to help you get the most out of your newly installed and configured FortiGate.

Best Practices

The Best Practices document is a collection of guidelines to ensure the most secure and reliable operation of FortiGates in a customer environment. It is updated periodically as new issues are identified.

This document can be found at <https://docs.fortinet.com/>.

The Fortinet Cookbook

The Fortinet Cookbook contains a variety of step-by-step examples of how to integrate a FortiGate into your network and apply features such as security profiles, wireless networking, and VPN.

Using the Cookbook, you can go from idea to execution in simple steps, configuring a secure network for better productivity with reduced risk.

The Fortinet Cookbook can be found at <http://cookbook.fortinet.com>.

The Fortinet Video Library

The Fortinet Video Library contains video tutorials showing how to configure various Fortinet products, including FortiGates. Many FortiGate videos are based on recipes from the FortiGate Cookbook.

The Fortinet Video Library can be found at <https://video.fortinet.com>. You can also [subscribe to Fortinet's YouTube channel](#).

The FortiOS Handbook

The FortiOS Handbook is the complete guide to FortiOS, covering a variety of FortiGate configurations. The Handbook is available as a single complete document online. Handbook chapters are also available as standalone documents.

The FortiOS Handbook can be found at <https://docs.fortinet.com/>.

Fortinet Support

You can also contact Fortinet Support for assistance. Read the following article found on the Fortinet Cookbook website: [How to work with Fortinet Support](#) to understand what type of support is available and to determine which level of support is right for you. For further information, go to <https://support.fortinet.com>.

Chapter 3 - Authentication

This Handbook chapter contains the following sections:

[Introduction to authentication](#) describes some basic elements and concepts of authentication.

[Authentication servers](#) describes external authentication servers, where a FortiGate unit fits into the topology, and how to configure a FortiGate unit to work with that type of authentication server.

[Users and user groups](#) describes the different types of user accounts and user groups. Authenticated access to resources is based on user identities and user group membership. Two-factor authentication methods, including FortiToken, provide additional security.

[Managing guest access](#) explains how to manage temporary accounts for visitors to your premises.

[Configuring authenticated access](#) provides detailed procedures for setting up authenticated access in security policies and authenticated access to VPNs.

[Captive portals](#) describes how to authenticate users through a web page that the FortiGate unit presents in response to any HTTP request until valid credentials are entered. This can be used for wired or WiFi network interfaces.

[Certificate-based authentication](#) describes authentication by means of X.509 certificates.

[Single sign-on using a FortiAuthenticator unit](#) describes how to use a FortiAuthenticator unit as an Single Sign-On (SSO) agent that can integrate with external network authentication systems such as RADIUS and LDAP to gather user logon information and send it to the FortiGate unit. Users can also log on through a FortiAuthenticator-based web portal or the FortiClient SSO Mobility Agent.

[Single sign-on to Windows AD](#) describes how to set up SSO in a Windows AD network by configuring the FortiGate unit to poll domain controllers for information user logons and user privileges.

[Agent-based FSSO](#) describes how to set up SSO in Windows AD, Citrix, or Novell networks by installing Fortinet Single Sign-On (FSSO) agents on domain controllers. The FortiGate unit receives information about user logons and allows access to network resources based on user group memberships.

[SSO using RADIUS accounting records](#) describes how to set up SSO in a network that uses RADIUS authentication. In this configuration, the RADIUS server send RADIUS accounting records to the FortiGate unit when users log on or off the network. The record includes a user group name that can be used in FortiGate security policies to determine which resources each user can access.

[Monitoring authenticated users](#) describes FortiOS authenticated user monitor screens.

[Examples and troubleshooting](#) provides configuration examples and troubleshooting suggestions.

What's new in FortiOS 6.0.1

The following list contains new authentication features added in FortiOS 6.0.1. Click on a link to navigate to that section for further information.

- [VMware Horizon support for TSSAgent \(for more information, see "Agent-based FSSO" on page 276\).](#)

What's new in FortiOS 6.0

The following list contains new authentication features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- ["Authentication binds to MAC address" on page 148](#)
- ["RADIUS accounting start message options" on page 160](#)
- ["UPN processing method and filter name" on page 167](#)
- ["IPv6 TACACS+ server IP address" on page 173](#)
- ["FortiToken extension to comply with PCI 3.2" on page 186](#)
- ["Login credentials for guest users shown in clear text on GUI and voucher" on page 196](#)
- ["Kerberos authentication for explicit web and transparent web proxy users" on page 218](#)
- ["IPv6 and captive portals" on page 235](#)
- ["Certificate-related protocols" on page 243](#)
- ["Agentless NTLM support" on page 285](#)
- ["Collector agent installation" on page 286](#)
- ["IPv6 support for FSSO" on page 315](#)
- [Quarantine CLI commands moved from switch-controller to user \(for more information, see the \[FortiOS 6.0 CLI Reference\]\(#\)\)](#)
- [New Amazon device category \(for more information, see the \[FortiOS 6.0 CLI Reference\]\(#\)\)](#)

Introduction to authentication

Identifying users and other computers—authentication—is a key part of network security. This section describes some basic elements and concepts of authentication.

The following topics are included in this section:

- [What is authentication?](#)
- [Methods of authentication](#)
- [Types of authentication](#)
- [User's view of authentication](#)
- [FortiGate administrator's view of authentication](#)

What is authentication?

Businesses need to authenticate people who have access to company resources. In the physical world this may be a swipe card to enter the building, or a code to enter a locked door. If a person has this swipe card or code, they have been authenticated as someone allowed in that building or room.

Authentication is the act of confirming the identity of a person or other entity. In the context of a private computer network, the identities of users or host computers must be established to ensure that only authorized parties can access the network. The FortiGate unit enables controlled network access and applies authentication to users of security policies and VPN clients.

Methods of authentication

FortiGate unit authentication is divided into three basic types: password authentication for people, certificate authentication for hosts or endpoints, and two-factor authentication for additional security beyond just passwords. An exception to this is that FortiGate units in an HA cluster and FortiManager units use password authentication.

Password authentication verifies individual user identities, but access to network resources is based on membership in user groups. For example, a security policy can be configured to permit access only to the members of one or more user groups. Any user who attempts to access the network through that policy is then authenticated through a request for their username and password.

Methods of authentication include:

- [Local password authentication](#)
- [Server-based password authentication](#)
- [Certificate-based authentication](#)
- [Two-factor authentication](#)

Local password authentication

The simplest authentication is based on user accounts stored locally on the FortiGate unit. For each account, a username and password is stored. The account also has a disable option so that you can suspend the account without deleting it.

Local user accounts work well for a single-FortiGate installation. If your network has multiple FortiGate units that will use the same accounts, the use of an external authentication server can simplify account configuration and maintenance.

You can create local user accounts in the web-based manager under **User & Device > User Definition**. This page is also used to create accounts where an external authentication server stores and verifies the password.

Server-based password authentication

Using external authentication servers is desirable when multiple FortiGate units need to authenticate the same users, or where the FortiGate unit is added to a network that already contains an authentication server. FortiOS supports the use of LDAP, RADIUS, TACACS+, AD, or POP3 servers for authentication.

When you use an external authentication server to authenticate users, the FortiGate unit sends the user's entered credentials to the external server. The password is encrypted. The server's response indicates whether the supplied credentials are valid or not.

You must configure the FortiGate unit to access the external authentication servers that you want to use. The configuration includes the parameters that authenticate the FortiGate unit to the authentication server.

You can use external authentication servers in two ways:

- Create user accounts on the FortiGate unit. However, instead of storing each user's password, specify the server used to authenticate that user. As with accounts that store the password locally, you add these users to appropriate user groups.
- Add the authentication server to user groups. Any user who has an account on the server can be authenticated and have the access privileges of the FortiGate user group. Optionally, when an LDAP server is a FortiGate user group member, you can limit access to users who belong to specific groups defined on the LDAP server.

Certificate-based authentication

An RSA X.509 server certificate is a small file issued by a certificate authority (CA) that is installed on a computer or FortiGate unit to authenticate itself to other devices on the network. When one party on a network presents the certificate as authentication, the other party can validate that the certificate was issued by the CA. The identification is therefore as trustworthy as the CA that issued the certificate.

To protect against compromised or misused certificates, CAs can revoke any certificate by adding it to a certificate revocation list (CRL). Certificate status can also be checked online using the Online Certificate Status Protocol (OCSP).

RSA X.509 certificates are based on public-key cryptography, in which there are two keys: the private key and the public key. Data encrypted with the private key can be decrypted only with the public key, and the other way round. As the names suggest, the private key is never revealed to anyone and the public key can be freely distributed. Encryption with the recipient's public key creates a message that only the intended recipient can read. Encryption with the sender's private key creates a message whose authenticity is proven because it can be decrypted only with the sender's public key.

Server certificates contain a signature string encrypted with the CA's private key. The CA's public key is contained in a CA root certificate. If the signature string can be decrypted with the CA's public key, the certificate is genuine.

Certificate authorities

A CA can be:

- an organization, such as VeriSign Inc., that provides certificate services
- a software application, such as Microsoft Certificate Services or OpenSSH

For a company web portal or customer-facing SSL VPN, a third-party certificate service has some advantages. The CA certificates are already included in popular web browsers and customers trust the third-party. On the other hand, third-party services have a cost.

For administrators and for employee VPN users, the local CA based on a software application provides the required security at low cost. You can generate and distribute certificates as needed. If an employee leaves the organization, you can simply revoke their certificate.

Certificates for users

FortiGate unit administrators and SSL VPN users can install certificates in their web browsers to authenticate themselves. If the FortiGate unit uses a CA-issued certificate to authenticate itself to the clients, the browser will also need the appropriate CA certificate.

FortiGate IPsec VPN users can install server and CA certificates according to the instructions for their IPsec VPN client software. The FortiClient Endpoint Security application, for example, can import and store the certificates required by VPN connections.

FortiGate units are also compatible with some Public Key Infrastructure systems. For an example of this type of system, see [RSA ACE \(SecurID\) servers on page 176](#).

Two-factor authentication

A user can be required to provide both something they know (their username and password combination) and something they have (certificate or a random token code). Certificates are installed on the user's computer.

Two-factor authentication is available for PKI users. For more information, see [Certificate on page 186](#).

Another type of two-factor authentication is to use a randomly generated token (multi-digit number) along with the username and password combination. One method is a FortiToken — a one-time password (OTP) generator that generates a unique code every 60 seconds. Others use email or SMS text messaging to deliver the random token code to the user or administrator.

When one of these methods is configured, the user enters this code at login after the username and password have been verified. The FortiGate unit verifies the token code after as well as the password and username. For more information, see [Two-factor authentication on page 185](#)

Authentication binds to MAC address

In previous FortiOS versions, firewall authentication was source IP based, thus there was no action in response to a MAC address change. This was a security flaw that allowed an unauthenticated user to access restricted resources, especially in a WiFi environment where the IP and MAC binding changed frequently.

MAC addresses can now be bound with the user identity so that the MAC address is matched while matching an auth logon.

Types of authentication

FortiOS supports two different types of authentication based on your situation and needs: security policy authentication and Virtual Private Network (VPN) authentication.

Security policy authentication is easily applied to all users logging on to a network, or network service. For example if a group of users on your network such as the accounting department who have access to sensitive data need to access the Internet, it is a good idea to make sure the user is a valid user and not someone trying to send company secrets to the Internet. Security policy authentication can be applied to as many or as few users as needed, and it supports a number of authentication protocols to easily fit with your existing network.

VPN authentication enables secure communication with hosts located outside the company network, making them part of the company network while the VPN tunnel is operating. Authentication applies to the devices at both ends of the VPN and optionally VPN users can be authenticated as well.

Security policy authentication

Security policies enable traffic to flow between networks. Optionally, the policy can allow access only to specific originating addresses, device types, users or user groups. Where access is controlled by user or user group, users must authenticate by entering valid username and password credentials.

The user's authentication expires if the connection is idle for too long, five minutes by default but that can be customized.

Security policies are the mechanism for FSSO, NTLM, certificate based, and RADIUS SSO authentication.

FSSO

Fortinet Single Sign on (FSSO) provides seamless authentication support for Microsoft Windows Active Directory (AD) and Novell eDirectory users in a FortiGate environment.

On a Microsoft Windows or Novell network, users authenticate with the Active Directory or Novell eDirectory at login. FSSO provides authentication information to the FortiGate unit so that users automatically get access to permitted resources. See [Introduction to agent-based FSSO on page 276](#).

NTLM

The NT LAN Manager (NTLM) protocol is used when the MS Windows Active Directory (AD) domain controller can not be contacted. NTLM is a browser-based method of authentication.

The FSSO software is installed on each AD server and the FortiGate unit is configured to communicate with each FSSO client. When a user successfully logs into their Windows PC (and is authenticated by the AD Server), the FSSO client communicates the user's name, IP address, and group login information to the FortiGate unit. The FortiGate unit sets up a temporary access policy for the user, so when they attempt access through the firewall they do not need to re-authenticate. This model works well in environments where the FSSO client can be installed on all AD servers.

In system configurations where it is not possible to install FSSO clients on all AD servers, the FortiGate unit must be able to query the AD servers to find out if a user has been properly authenticated. This is achieved using the NTLM messaging features of Active Directory and Internet Explorer.

Even when NTLM authentication is used, the user is not asked again for their username and password. Internet Explorer stores the user's credentials and the FortiGate unit uses NTLM messaging to validate them in the Windows AD environment.

Note that if the authentication reaches the timeout period, the NTLM message exchange restarts. For more information on NTLM, see [NTLM authentication on page 217](#) and [FSSO NTLM authentication support on page 282](#).

Certificates

Certificates can be used as part of a policy. All users being authenticated against the policy are required to have the proper certificate. See [Certificate-based authentication on page 241](#)

RADIUS SSO

RADIUS Single Sign-On (RSSO) is a remote authentication method that does not require any local users to be configured, and relies on RADIUS Start records to provide the FortiGate unit with authentication information. That information identifies the user and user group, which is then matched using a security policy. See [SSO using RADIUS accounting records on page 321](#).

FortiGuard web filter override authentication

Optionally, users can be allowed the privilege of overriding FortiGuard Web Filtering to view blocked web sites. Depending on the override settings, the override can apply to the user who requested it, the entire user group to which the user belongs, or all users who share the same web filter profile. As with other FortiGate features, access to FortiGuard overrides is controlled through user groups. Firewall and Directory Services user groups are eligible for the override privilege. For more information about web filtering and overrides, see the UTM chapter of this FortiOS Handbook.

VPN authentication

Authentication involves authenticating the user. In IPsec VPNs authenticating the user is optional, but authentication of the peer device is required.

This section includes:

- [Authenticating IPsec VPN peers \(devices\)](#)
- [Authenticating IPsec VPN users](#)
- [Authenticating SSL VPN users](#)
- [Authenticating PPTP and L2TP VPN users](#)

Authenticating IPsec VPN peers (devices)

A VPN tunnel has one end on a local trusted network, and the other end is at a remote location. The remote peer (device) must be authenticated to be able to trust the VPN tunnel. Without that authentication, it is possible for a malicious hacker to masquerade as a valid VPN tunnel device and gain access to the trusted local network.

The three ways to authenticate VPN peers are with a pre-shared key, RSA X.509 certificate, or a specific peer ID value.

The simplest way for IPsec VPN peers to authenticate each other is through the use of a pre-shared key, also called a shared secret. The pre-shared key is a text string used to encrypt the data exchanges that establish the VPN tunnel. The pre-shared key must be six or more characters. The VPN tunnel cannot be established if the two peers do not use the same key. The disadvantage of pre-shared key authentication is that it can be difficult to securely distribute and update the pre-shared keys.

RSA X.509 certificates are a better way for VPN peers to authenticate each other. Each peer offers a certificate signed by a Certificate Authority (CA) which the other peer can validate with the appropriate CA root certificate. For more information about certificates, see [Certificate-based authentication on page 241](#).

You can supplement either pre-shared key or certificate authentication by requiring the other peer to provide a specific peer ID value. The peer ID is a text string configured on the peer device. On a FortiGate peer or FortiClient Endpoint Security peer, the peer ID provided to the remote peer is called the Local ID.

Authenticating IPsec VPN users

An IPsec VPN can be configured to accept connections from multiple dynamically addressed peers. You would do this to enable employees to connect to the corporate network while traveling or from home. On a FortiGate unit, you create this configuration by setting the **Remote Gateway** to **Dialup User**.

It is possible to have an IPsec VPN in which remote peer devices authenticate using a common pre-shared key or a certificate, but there is no attempt to identify the user at the remote peer. To add user authentication, you can do one of the following:

- require a unique pre-shared key for each peer
- require a unique peer ID for each peer
- require a unique peer certificate for each peer
- require additional user authentication (XAuth)

The peer ID is a text string configured on the peer device. On a FortiGate peer or FortiClient Endpoint Security peer, the peer ID provided to the remote peer is called the Local ID.

Authenticating SSL VPN users

SSL VPN users can be

- user accounts with passwords stored on the FortiGate unit
- user accounts authenticated by an external RADIUS, LDAP or TACACS+ server
- PKI users authenticated by certificate

You need to create a user group for your SSL VPN. Simply create a firewall user group, enable SSL VPN access for the group, and select the web portal the users will access.

SSL VPN access requires an SSL VPN security policy that permits access to members of your user group.

Authenticating PPTP and L2TP VPN users

PPTP and L2TP are older VPN tunneling protocols that do not provide authentication themselves. FortiGate units restrict PPTP and L2TP access to users who belong to one specified user group. Users authenticate themselves to the FortiGate unit by username/password. You can configure PPTP and L2TP VPNs only in the CLI. Before you configure the VPN, create a firewall user group and add to it the users who are permitted to use the VPN. Users are authenticated when they attempt to connect to the VPN. For more information about configuring PPTP or L2TP VPNs, see the FortiGate CLI Reference.

Single sign-on authentication for users

Single sign-on (SSO) means that users logged on to a computer network are authenticated for access to network resources through the FortiGate unit without having to enter their username and password again. FortiGate units directly provide SSO capability for:

- Microsoft Windows networks using either Active Directory or NTLM authentication
- Novell networks, using eDirectory

In combination with a FortiAuthenticator unit, the FortiGate unit can provide SSO capability that integrates multiple external network authentication systems such as Windows Active Directory, Novell e-Directory, RADIUS and LDAP. The FortiAuthenticator unit gathers user logon information from all of these sources and sends it to the FortiGate unit.

Through the SSO feature, the FortiGate unit knows the username, IP address, and external user groups to which the user belongs. When the user tries to access network resources, the FortiGate unit selects the appropriate security policy for the destination. If the user belongs to one of the permitted user groups, the connection is allowed.

For detailed information about SSO, see

- [Single sign-on using a FortiAuthenticator unit on page 263](#)
- [Agent-based FSSO on page 276](#)

User's view of authentication

From the user's point of view, they see a request for authentication when they try to access a protected resource, such as an FTP repository of intellectual property or simply access a website on the Internet. The way the request is presented to the user depends on the method of access to that resource.

VPN authentication usually controls remote access to a private network.

Web-based user authentication

Security policies usually control browsing access to an external network that provides connection to the Internet. In this case, the FortiGate unit requests authentication through the web browser.

The user types a username and password and then selects **Continue** or **Login**. If the credentials are incorrect, the authentication screen is redisplayed with blank fields so that the user can try again. When the user enters valid credentials, access is granted to the required resource. In some cases, if a user tries to authenticate several times without success, a message appears, such as: "Too many bad login attempts. Please try again in a few minutes." This indicates the user is locked out for a period of time. This prevents automated brute force password hacking attempts. The administrator can customize these settings if required.



After a defined period of user inactivity (the authentication timeout, defined by the FortiGate administrator), the user's access expires. The default is 5 minutes. To access the resource, the user will have to authenticate again.

VPN client-based authentication

A VPN provides remote clients with access to a private network for a variety of services that include web browsing, email, and file sharing. A client program such as FortiClient negotiates the connection to the VPN and manages the user authentication challenge from the FortiGate unit.

FortiClient can store the username and password for a VPN as part of the configuration for the VPN connection and pass them to the FortiGate unit as needed. Or, FortiClient can request the username and password from the user when the FortiGate unit requests them.



Social ID data from FortiClient is recorded so that if an email or phone number is changed on FortiClient, the new values are updated on the FortiGate.

The data is sent in KeepAlive messages in the following format:

```
USR_NAME|<full name for the service account>|USR_
EMAIL|<email for the service
account>|SERVICE|<os|custom|linkedin|google|salesforce>|
```

SSL VPN is a form of VPN that can be used with a standard Web browser. There are two modes of SSL VPN operation (supported in NAT/Route mode only):

- web-only mode, for remote clients equipped with a web-browser only
- tunnel mode, for remote computers that run a variety of client and server applications.



After a defined period of user inactivity on the VPN connection (the idle timeout, defined by the FortiGate administrator), the user's access expires. The default is 30 minutes. To access the resource, the user will have to authenticate again.

FortiGate administrator's view of authentication

Authentication is based on user groups. The FortiGate administrator configures authentication for security policies and VPN tunnels by specifying the user groups whose members can use the resource. Some planning is required to determine how many different user groups need to be created. Individual user accounts can belong to multiple groups, making allocation of user privileges very flexible.

A member of a user group can be:

- a user whose username and password are stored on the FortiGate unit
- a user whose name is stored on the FortiGate unit and whose password is stored on a remote or external authentication server
- a remote or external authentication server with a database that contains the username and password of each person who is permitted access

The general process of setting up authentication is as follows:

1. If remote or external authentication is needed, configure the required servers.
2. Configure local and peer (PKI) user identities. For each local user, you can choose whether the FortiGate unit or a remote authentication server verifies the password. Peer members can be included in user groups for use in security policies.
3. Create user groups.
4. Add local/peer user members to each user group as appropriate. You can also add an authentication server to a user group. In this case, all users in the server's database can authenticate. You can only configure peer user groups through the CLI.
5. Configure security policies and VPN tunnels that require authenticated access.

For authentication troubleshooting, see the specific chapter for the topic or for general issues see [Troubleshooting on page 348](#).

General authentication settings

Go to **User & Device > Authentication Settings** to configure authentication timeout, protocol support, and authentication certificates.

When user authentication is enabled within a security policy, the authentication challenge is normally issued for any of the four protocols (depending on the connection protocol):

- HTTP (can also be set to redirect to HTTPS)
- HTTPS
- FTP
- Telnet

The selections made in the **Protocol Support** list of **Authentication Settings** control which protocols support the authentication challenge. Users must connect with a supported protocol first so they can subsequently connect with other protocols. If HTTPS is selected as a method of protocol support, it allows the user to authenticate with a customized Local certificate.

When you enable user authentication within a security policy, the security policy user will be challenged to authenticate. For user ID and password authentication, users must provide their user names and passwords. For certificate authentication (HTTPS or HTTP redirected to HTTPS only), you can install customized certificates on the unit and the users can also have customized certificates installed on their browsers. Otherwise, users will see a warning message and have to accept a default Fortinet certificate.

Authentication Timeout	Enter a length of time in minutes, from 1 to 4320 (72 hours). Authentication timeout controls how long an authenticated firewall connection can be idle before the user must authenticate again. The default value is 5.
Protocol Support	Select the protocols to challenge during firewall user authentication.
Certificate	If using HTTPS protocol support, select the local certificate to use for authentication. Available only if HTTPS protocol support is selected.
Apply	Select to apply the selections for user authentication settings.

Authentication servers

FortiGate units support the use of external authentication servers. An authentication server can provide password checking for selected FortiGate users or it can be added as a member of a FortiGate user group.

If you are going to use authentication servers, you must configure the servers before you configure FortiGate users or user groups that require them.



Mac OS and iOS devices, including iPhones and iPads, can perform user authentication with FortiOS units using RADIUS servers, but not with LDAP or TACACS+ servers.

This section includes the following topics:

- [FortiAuthenticator servers](#)
- [RADIUS servers](#)
- [LDAP servers](#)
- [TACACS+ servers](#)
- [POP3 servers](#)
- [SSO servers](#)
- [RSA ACE \(SecurID\) servers](#)

FortiAuthenticator servers

FortiAuthenticator is an Authentication, Authorization, and Accounting (AAA) server, that includes a RADIUS server, an LDAP server, and can replace the FSSO Collector Agent on a Windows AD network. Multiple FortiGate units can use a single FortiAuthenticator for FSSO, remote authentication, and FortiToken management.

For more information, see the [FortiAuthenticator Administration Guide](#).

RADIUS servers

Remote Authentication and Dial-in User Service (RADIUS) is a broadly supported client-server protocol that provides centralized authentication, authorization, and accounting functions. RADIUS clients are built into gateways that allow access to networks such as Virtual Private Network servers, Network Access Servers (NAS), as well as network switches and firewalls that use authentication. FortiGate units fall into the last category.

RADIUS servers use UDP packets to communicate with the RADIUS clients on the network to authenticate users before allowing them access to the network, to authorize access to resources by appropriate users, and to account or bill for those resources that are used. RADIUS servers are currently defined by RFC 2865 (RADIUS) and RFC 2866 (Accounting), and listen on either UDP ports 1812 (authentication) and 1813 (accounting) or ports 1645 (authentication) and 1646 (accounting) requests. RADIUS servers exist for all major operating systems.

You must configure the RADIUS server to accept the FortiGate unit as a client. FortiGate units use the authentication and accounting functions of the RADIUS server.



FortiOS does not accept all characters from auto generated keys from MS Windows 2008. These keys are very long and as a result RADIUS authentication will not work. Maximum key length for MS Windows 2008 is 128 bytes. In older versions of FSAE, it was 40 bytes.

Microsoft RADIUS servers

Microsoft Windows Server 2000, 2003, and 2008 have RADIUS support built-in. Microsoft specific RADIUS features are defined in RFC 2548. The Microsoft RADIUS implementation can use Active Directory for user credentials.

For details on Microsoft RADIUS server configurations, refer to Microsoft documentation.

RADIUS user database

The RADIUS user database is commonly an SQL or LDAP database, but can also be any combination of:

- usernames and passwords defined in a configuration file
- user account names and passwords configured on the computer where the RADIUS server is installed.

If users are members of multiple RADIUS groups, then the user group authentication timeout value does not apply. See [Membership in multiple groups on page 199](#).

RADIUS authentication with a FortiGate unit

To use RADIUS authentication with a FortiGate unit

- configure one or more RADIUS servers on the FortiGate unit
- assign users to a RADIUS server

When a configured user attempts to access the network, the FortiGate unit will forward the authentication request to the RADIUS server which will match the username and password remotely. Once authenticated the RADIUS server passes the authorization granted message to the FortiGate unit which grants the user permission to access the network.

The RADIUS server uses a “shared secret” key along with MD5 hashing to encrypt information passed between RADIUS servers and clients, including the FortiGate unit. Typically only user credentials are encrypted. Additional security can be configured through IPsec tunnels by placing the RADIUS server behind another VPN gateway.

RADIUS attribute value pairs

RADIUS packets include a set of attribute value pairs (AVP) to identify information about the user, their location and other information. The FortiGate unit sends the following RADIUS attributes.

FortiOS supported RADIUS attributes

RADIUS Attribute	Name	Description	AVP type
1	Acct-Session-ID	Unique number assigned to each start and stop record to make it easy to match them, and to eliminate duplicate records.	44

RADIUS Attribute	Name	Description	AVP type
2	Username	Name of the user being authenticated	1
3	NAS-Identifier	Identifier or IP address of the Network Access Server (NAS) that is requesting authentication. In this case, the NAS is the FortiGate unit.	32
4	Framed-IP-Address	Address to be configured for the user.	8
5	Fortinet-VSA	See Vendor-specific attributes on page 158	26
6	Acct-Input-Octets	Number of octets received from the port over the course of this service being provided. Used to charge the user for the amount of traffic they used.	42
7	Acct-Output-Octets	Number of octets sent to the port while delivering this service. Used to charge the user for the amount of traffic they used.	43
8	NAS-IP-Address	IP address of the Network Access Server (NAS) that is requesting authentication. In this case, the NAS is the FortiGate unit.	4
9	Called-Station-Id	Used to send the telephone number the user called as part of the Access-Request packet.	30
10	Framed-IP-Address	IP address to be configured for the user, by sending the IP address of a user to the RADIUS server in the Access-Request packet.	8
11	Event-Timestamp	Records the time that the event occurred on the NAS. The timestamp is measured in seconds since January 1, 1970 00:00 UTC. Before the Event-Timestamp attribute can be sent in a packet, make sure that the correct time is set on the FortiGate.	55
12	Class	Used in accounting packets and requests for firewall, WiFi, and proxy authentication. The attribute is returned in Access-Account message and is added to all accounting packets.	25

The following table describes the supported authentication events and the RADIUS attributes that are sent in the RADIUS accounting message.

RADIUS attributes sent in RADIUS accounting message

Authentication Method	RADIUS Attributes						
	1	2	3	4	5	6	7
Web	X	X	X		X		
XAuth of IPsec (without DHCP)	X	X	X		X		
XAuth of IPsec (with DHCP)	X	X	X	X	X		
PPTP/L2TP (in PPP)	X	X	X	X	X	X	X
SSL-VPN	X	X	X	X	X		

External captive portal POST message

In external RADIUS captive portal, the captive portal web page is a script that gathers the user's logon credentials and sends it back to the FortiGate as a POST message. Session URL parameters are sent from the client in a POST messages, and in the redirect. These parameters are separated by **&** characters (see examples below):

POST message to redirect server:

```
http://<redirectserver>/index2.php/?login&post=http://192.168.200.1:1000/fgtauth&magic=02050f889bc21644&usermac=54:26:96:16:a2:45&apmac=00:09:0f:b9:f4:c0&apip=127.0.0.1&userip=192.168.200.2
```

POST message back to the FortiGate:

```
http://FGT_IP_addr:1000/fgtauth
```

The magic text data, provided in the initial FortiGate request to the web server, contains the username, password parameters:

```
magic=00050c839182f095&username=<username>&password=<password>
```

Vendor-specific attributes

Vendor specific attributes (VSA) are the method RADIUS servers and client companies use to extend the basic functionality of RADIUS. Some major vendors, such as Microsoft, have published their VSAs, however many do not.

In order to support vendor-specific attributes (VSA), the RADIUS server requires a dictionary to define which VSAs to support. This dictionary is typically supplied by the client or server vendor.

The Fortinet RADIUS vendor ID is 12356.

The FortiGate unit RADIUS VSA dictionary is supplied by Fortinet and is available through the Fortinet Knowledge Base (<http://kb.forticare.com>) or through Technical Support. Fortinet's dictionary for FortiOS 4.0 and up is configured this way:

```
##
```

```

Fortinet's VSA's
#
VENDOR fortinet 12356
BEGIN-VENDOR fortinet
ATTRIBUTE Fortinet-Group-Name 1 string
ATTRIBUTE Fortinet-Client-IP-Address 2 ipaddr
ATTRIBUTE Fortinet-Vdom-Name 3 string
ATTRIBUTE Fortinet-Client-IPv6-Address 4 octets
ATTRIBUTE Fortinet-Interface-Name 5 string
ATTRIBUTE Fortinet-Access-Profile 6 string
#
# Integer Translations
#
END-VENDOR Fortinet

```

Note that using the Fortinet-Vdom-Name, users can be tied to a specific VDOM on the FortiGate unit. See the documentation provided with your RADIUS server for configuration details.

RADIUS CoA support

As of FortiOS 5.4, RADIUS Change of Authorization (CoA) settings can be configured via the CLI. CoA is a common feature in user authentication that provides the ability to change authentication attributes for sessions even after they have authenticated.

User, user group, and captive portal authentication supports RADIUS CoA, when the back end authentication server is RADIUS. The main use case of this feature is with external captive portal, where it can be used to disconnect hotspot users when their time, credit, or bandwidth has been used up.

The commands below control CoA settings.

1. Set the name of the FortiAP connected to the FortiGate as a location identifier.

```

config system global
set alias <name>

```

2. Set URL of external authentication logout server.

```

config vdom
edit root
config wireless-controller vap
edit <example>
set security captive-portal
set external-logout

```

3. Set URL of external authentication logout server

```

config vdom
edit root
config system interface
edit <example>
set security captive-portal
set security-external-logout

```

4. Set class name(s) included in an Access-Accept message.

```

config vdom
edit root
config user radius
edit accounting
set class <"A1=aaa" "B2=bbb" "C3=ccc">

```

Role-based access control

In role-based access control (RBAC), network administrators and users have varying levels of access to network resources based on their role, and that role's requirement for access specific resources. For example, a junior accountant does not require access to the sales presentations, or network user account information.

There are three main parts to RBAC: role assignment, role authorization, and transaction authorization. Role assignment is accomplished when someone in an organization is assigned a specific role by a manager or HR. Role authorization is accomplished when a network administrator creates that user's RADIUS account and assigns them to the required groups for that role. Transaction authorization occurs when that user logs on and authenticates before performing a task.

RBAC is enforced when FortiOS network users are remotely authenticated via a RADIUS server. For users to authenticate, a security policy must be matched. That policy only matches a specific group of users. If VDOMs are enabled, the matched group will be limited to a specific VDOM. Using this method network administrators can separate users into groups that match resources, protocols, or VDOMs. It is even possible to limit users to specific FortiGate units if the RADIUS servers serve multiple FortiOS units.

For more information on security policies, see [Authentication in security policies on page 212](#).

RADIUS password encoding

Certain RADIUS servers use ISO-8859-1 password encoding instead of others such as UTF-8. In these instances, the server will fail to authenticate the user, if the user's password is using UTF-8.

CLI syntax

```
config user radius
  edit <example>
    set password-encoding <auto | ISO-8859-1>
  end
```

This option will be skipped if the `auth-type` is neither `auto` nor `pap`.

RADIUS accounting start message options

Administrators can now choose between sending accounting start messages to all configured accounting servers, or just the one server that was previously connected.

Syntax

```
config user radius
  edit <name>
    set acct-interim-interval <seconds>
    set acct-all-servers {enable | disable}
  next
end
```

Configuring the FortiGate unit to use a RADIUS server

The information you need to configure the FortiGate unit to use a RADIUS server includes

- the RADIUS server's domain name or IP address
- the RADIUS server's shared secret key.

You can optionally specify the NAS IP or Called Station ID. When configuring the FortiGate to use a RADIUS server, the FortiGate is a Network Access Server (NAS). If the FortiGate interface has multiple IP addresses, or you want the RADIUS requests to come from a different address you can specify it here. Called Station ID applies to carrier networks. However, if the NAS IP is not included in the RADIUS configuration, the IP of the FortiGate unit interface that communicates with the RADIUS server is used instead.

A maximum of 10 remote RADIUS servers can be configured on the FortiGate unit. One or more servers must be configured on FortiGate before remote users can be configured. To configure remote users, see [Local and remote users on page 182](#).

On the FortiGate unit, the default port for RADIUS traffic is 1812. Some RADIUS servers use port 1645. If this is the case with your server, you can either:

- Re-configure the RADIUS server to use port 1812. See your RADIUS server documentation for more information on this procedure.

or

- Change the FortiGate unit default RADIUS port to 1645 using the CLI:

```
config system global
    set radius-port 1645
end
```

One wildcard admin account can be added to the FortiGate unit when using RADIUS authentication. This uses the wildcard character to allow multiple admin accounts on RADIUS to use a single account on the FortiGate unit. See [Example — wildcard admin accounts - CLI on page 168](#).

To configure the FortiGate unit for RADIUS authentication - web-based manager:

1. Go to **User & Device > RADIUS Servers** and select **Create New**.
2. Enter the following information and select **OK**.

Name	A name to identify the RADIUS server on the FortiGate unit.
Authentication method	If you know the RADIUS server uses a specific authentication protocol, select Specify and select it from the list. Otherwise select Default . The Default option will usually work.
NAS IP	Enter the IP address to be used as the NAS-IP-Address attribute in RADIUS access requests.
Include in every user group	When enabled this RADIUS server will automatically be included in all user groups. This is useful if all users will be authenticating with the remote RADIUS server.

Primary Server

IP/Name: Enter the domain name (such as fgt.exmaple.com) or the IP address of the RADIUS server.

Secret: Enter the server secret key, such as radiusSecret. This can be a maximum of 16 characters long. This must match the secret on the RADIUS primary server.

Test Connectivity: Test the validity of the **IP/Name**.

Test User Connectivity: Test the validity of the **Secret**.

Secondary Server

IP/Name: Optionally enter the domain name (such as fgt.exmaple.com) or the IP address of the secondary RADIUS server.

Secret: Optionally, enter the secondary server secret key, such as radiusSecret2. This can be a maximum of 16 characters long. This must match the secret on the RADIUS secondary server.

Test Connectivity: Test the validity of the **IP/Name**.

Test User Connectivity: Test the validity of the **Secret**.



For MAC OS and iOS devices to authenticate, you must use MS-CHAP-v2 authentication. In the CLI, the command is `set auth-type ms_chap_v2`.

3. Select **OK**.

To configure the FortiGate unit for RADIUS authentication - CLI example:

```
config user radius
  edit ourRADIUS
    set auth-type auto
    set server 10.11.102.100
    set secret radiusSecret
  end
```

For more information about RADIUS server options, refer to the FortiGate CLI Reference.

Troubleshooting RADIUS

To test the connection to the RADIUS server use the following command:

```
diagnose test authserver radius-direct <server_name or IP> <port number> <secret>
```

For the port number, enter -1 to use the default port. Otherwise enter the port number to check.

Test results show RADIUS server reachability, NAS client rejection, and invalid User/Password. Test also shows RADIUS Attributes returned from the RADIUS server.

LDAP servers

Lightweight Directory Access Protocol (LDAP) is an Internet protocol used to maintain authentication data that may include departments, people, groups of people, passwords, email addresses, and printers. LDAP consists of a data-representation scheme, a set of defined operations, and a request/response network.

The scale of LDAP servers range from big public servers such as BigFoot and Infospace, to large organizational servers at universities and corporations, to small LDAP servers for workgroups that may be using OpenLDAP. This document focuses on the institutional and workgroup applications of LDAP.

This section includes:

- [Components and topology](#)
- [LDAP directory organization](#)
- [Configuring the FortiGate unit to use an LDAP server](#)
- [Example — wildcard admin accounts - CLI](#)
- [Example of LDAP to allow dial-in through member-attribute - CLI](#)
- [Troubleshooting LDAP](#)

Components and topology

LDAP organization starts with directories. A directory is a set of objects with similar attributes organized in a logical and hierarchical way. Generally, an LDAP directory tree reflects geographic and organizational boundaries, with the Domain name system (DNS) names to structure the top level of the hierarchy. The common name identifier for most LDAP servers is `cn`, however some servers use other common name identifiers such as `uid`.

When LDAP is configured and a user is required to authenticate the general steps are:

1. The FortiGate unit contacts the LDAP server for authentication.
2. To authenticate with the FortiGate unit, the user enters a username and password.
3. The FortiGate unit sends this username and password to the LDAP server.
4. If the LDAP server can authenticate the user, the user is successfully authenticated with the FortiGate unit.
5. If the LDAP server cannot authenticate the user, the connection is refused by the FortiGate unit.

Binding

Binding is the step where the LDAP server authenticates the user. If the user is successfully authenticated, binding allows the user access to the LDAP server based on that user's permissions.

The FortiGate unit can be configured to use one of three types of binding:

- anonymous - bind using anonymous user search
- regular - bind using username/password and then search
- simple - bind using a simple password authentication without a search

You can use simple authentication if the user records all fall under one domain name (`dn`). If the users are under more than one `dn`, use the anonymous or regular type, which can search the entire LDAP database for the required username.

If your LDAP server requires authentication to perform searches, use the regular type and provide values for username and password.

Supported versions

The FortiGate unit supports LDAP protocol functionality defined in RFC 2251: Lightweight Directory Access Protocol v3, for looking up and validating user names and passwords. FortiGate LDAP supports all LDAP servers compliant with LDAP v3, including FortiAuthenticator. In addition, FortiGate LDAP supports LDAP over SSL/TLS, which can be configured only in the CLI.

FortiGate LDAP does not support proprietary functionality, such as notification of password expiration, which is available from some LDAP servers. FortiGate LDAP does not supply information to the user about why authentication failed.



LDAP user authentication is supported for PPTP, L2TP, IPsec VPN, and firewall authentication.

However, with PPTP, L2TP, and IPsec VPN, PAP (Packet Authentication Protocol) is supported, while CHAP (Challenge Handshake Authentication Protocol) is not.

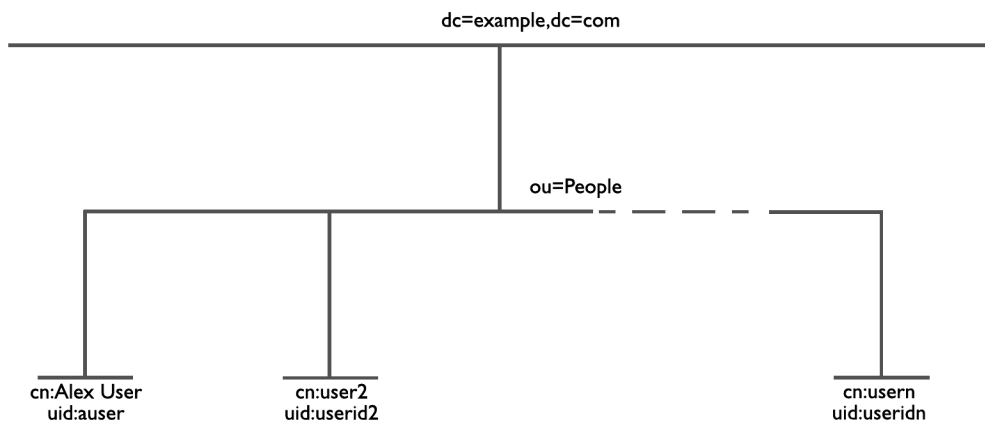
LDAP directory organization

To configure your FortiGate unit to work with an LDAP server, you need to understand the organization of the information on the server.

The top of the hierarchy is the organization itself. Usually this is defined as Domain Component (DC), a DNS domain. If the name contains a dot, such as `example.com`, it is written as two parts separated by a comma: `dc=example,dc=com`.

In this example, Common Name (CN) identifiers reside at the Organization Unit (OU) level, just below DC. The Distinguished Name (DN) is `ou=People,dc=example,dc=com`.

LDAP object hierarchy



In addition to the DN, the FortiGate unit needs an identifier for the individual person. Although the FortiGate unit GUI calls this the Common Name (CN), the identifier you use is not necessarily CN. On some servers, CN is the full name of a person. It might be more convenient to use the same identifier used on the local computer network. In this example, User ID (UID) is used.

Locating your identifier in the hierarchy

You need to determine the levels of the hierarchy from the top to the level that contain the identifier you want to use. This defines the DN that the FortiGate unit uses to search the LDAP database. Frequently used distinguished name elements include:

- uid (user identification)
- pw (password)
- cn (common name)
- ou (organizational unit)
- o (organization)
- c (country)

One way to test this is with a text-based LDAP client program. For example, OpenLDAP includes a client, `ldapsearch`, that you can use for this purpose.

Enter the following at the command line:

```
ldapsearch -x '(objectclass=*)'
```

The output is lengthy, but the information you need is in the first few lines:

```
version: 2
#
# filter: (objectclass=*)
# requesting: ALL

dn: dc=example,dc=com
dc: example
objectClass: top
objectClass: domain

dn: ou=People,dc=example,dc=com
ou: People
objectClass: top
objectClass: organizationalUnit
...
dn: uid=tbrown,ou=People,dc=example,dc=com
uid: tbrown
cn: Tom Brown
```

In the output above, you can see `tbrown` (uid) and `Tom Brown` (cn). Also note the dn is `ou=People, dc=example, dc=com`.

Configuring the FortiGate unit to use an LDAP server

After you determine the common name and distinguished name identifiers and the domain name or IP address of the LDAP server, you can configure the server on the FortiGate unit. The maximum number of remote LDAP servers that can be configured is 10.

One or more servers must be configured on FortiGate before remote users can be configured. To configure remote users, see [Local and remote users on page 182](#).

To configure the FortiGate unit for LDAP authentication - web-based manager:

1. Go to **User & Device > LDAP Servers** and select **Create New**.
2. Enter a **Name** for the LDAP server.
3. In **Server Name/IP** enter the server's FQDN or IP address.
4. If necessary, change the **Server Port** number. The default is port 389.
5. Enter the **Common Name Identifier** (20 characters maximum).
`cn` is the default, and is used by most LDAP servers.
6. In the **Distinguished Name** field, enter the base distinguished name for the server using the correct X.500 or LDAP format.
The FortiGate unit passes this distinguished name unchanged to the server. The maximum number of characters is 512.
If you don't know the distinguished name, leave the field blank and select the Query icon to the right of the field.
See [Using the query icon on page 167](#).
7. In the **Distinguished Name** field, enter the base distinguished name for the server using the correct X.500 or LDAP format.
The FortiGate unit passes this distinguished name unchanged to the server. The maximum number of characters is 512.
If you don't know the distinguished name, leave the field blank and select the Query icon to the right of the field.
See [Using the query icon on page 167](#).
8. In **Bind Type**, select **Regular**.
9. In **User DN**, enter the LDAP administrator's distinguished name.
10. In **Password**, enter the LDAP administrator's password.
11. Select **OK**.



To verify your Distinguished Name field is correct, you can select the **Test** button. If your DN field entry is valid, you will see the part of the LDAP database it defines. If your DN field entry is not valid, it will display an error message and return no information.

For detailed information about configuration options for LDAP servers, see the Online Help on your FortiGate unit or the FortiGate CLI Reference.

To configure the FortiGate unit for LDAP authentication - CLI example:

```
config user ldap
  edit ourLDAPsrv
    set server 10.11.101.160
    set cnid cn
    set dn cn=users,dc=office,dc=example,dc=com
    set type regular
    set username cn=administrator,cn=users,dc=office,dc=example,dc=com
    set password w5AiGVMLkgyPQ
    set password-expiry-warning enable
    set password-renewal enable
  end
```

password-expiry-warning and password-renewal

In SSLVPN, when an LDAP user is connecting to the LDAP server it is possible for them to receive any pending password expiry or renewal warnings. When the password renewal or expiry warning exists, SSLVPN users will see a prompt allowing them to change their password.

`password-expiry-warning` allows FortiOS to detect from the OpenLDAP server when a password is expiring or has expired using server controls or error codes. Please note that this is currently not supported for Windows AD LDAP.

`password-renewal` allows FortiOS to perform the online LDAP password renewal operations the LDAP server expects.

On an OpenLDAP server, when a user attempts to logon with an expired password they are allowed to logon but only to change their password.

When changing passwords on a Windows AD system, the connection must be SSL-protected.

UPN processing method and filter name

The following CLI commands available under `config user ldap` allow you to keep or strip the domain string of userPrincipalName (UPN) in the token as well as the search name for this kind of UPN.

Principle name peer user mode was only supported with Windows AD, and `fnbamd` has a hard coded `UserAccountControl` when doing LDAP authentication. `UserAccountControl` is a Windows AD specific attribute which doesn't exist in FortiAuthenticator or OpenLDAP.

To make LDAP query more flexible with different LDAP server versions, `account-key-filter` has been introduced to replace `account-key-name`. Enter the UPN attribute as the filter.

CLI syntax:

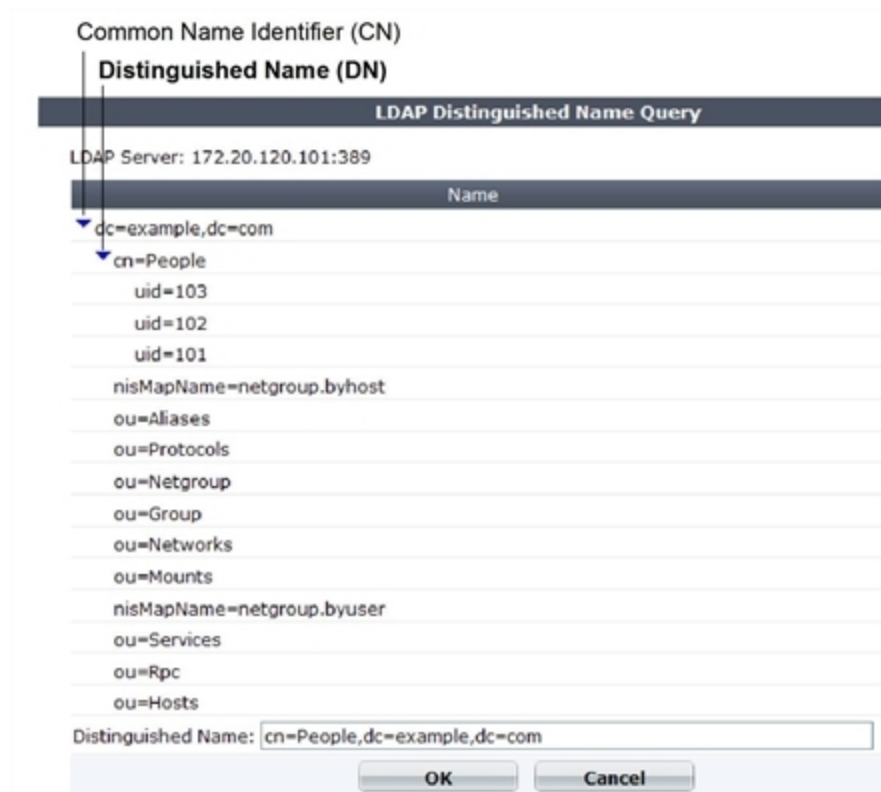
```
config user ldap
  set account-key-processing
  set account-key-filter
end
```

Using the query icon

The LDAP Distinguished Name Query list displays the LDAP directory tree for the LDAP server connected to the FortiGate unit. This helps you to determine the appropriate entry for the DN field. To see the distinguished name associated with the Common Name identifier, select the Expand icon next to the CN identifier. Select the DN from the list. The DN you select is displayed in the Distinguished Name field. Select OK and the Distinguished Name you selected will be saved in the Distinguished Name field of the LDAP Server configuration.

To see the users within the LDAP Server user group for the selected Distinguished Name, expand the Distinguished Name in the LDAP Distinguished Name Query tree.

LDAP server Distinguished Name Query tree



Non-blocking LDAP authentication

To support non-blocking LDAP authentication, fnbamd will create its own event-driven connection with LDAP servers over LDAP/LDAPS/STARTTLS, make it non-blocking, do CRL checking if necessary, and compose all LDAP requests using liblber (including bind, unbind, search, password renewal, password query, send request and receive response, and parse response). The whole process is done in one connection.

This doesn't change any openLDAP implementation but moves some data structure definitions and API definitions from some internal header files to public header files.

Example — wildcard admin accounts - CLI

A wildcard admin account is an administrator account with the wildcard option enabled. This option allows multiple different remote administration accounts to match one local administration account, avoiding the need to set up individual admin accounts on the FortiGate unit. Instead multiple LDAP admin accounts will all be able to use one FortiGate admin account.

The initial benefit of wildcard admin accounts is fast configuration of the FortiGate unit's administration account to work with your LDAP network. The many to one ratio saves on effort, and potential errors.

The ongoing benefit is that as long as the users on the LDAP system belong to that group, and the test admin user settings don't change on the FortiGate unit, no other work is required. This point is important as it can help avoid system updates or changes that would otherwise require changes to the LDAP administrator account

configuration. Even if a user is added to or removed from the LDAP group, no changes are required on the FortiGate unit.

Two potential issues with wildcard admin accounts are that multiple users may be logged on to the same account at the same time. This becomes an issue if they are changing the same information at the same time. The other potential issue is that security is reduced because multiple people have login access for the same account. If each user was assigned their own account, a hijacking of one account would not affect the other users.

Note that wildcard admin configuration also applies to RADIUS. When configuring for RADIUS, configure the RADIUS server, and RADIUS user group instead of LDAP. When using web-based management, wildcard admin is the only type of remote administrator account that does not require you to enter a password on account creation. That password is normally used when the remote authentication server is unavailable during authentication.

In this example, default values are used where possible. If a specific value is not mentioned, it is set to its default value.

Configuring the LDAP server

The important parts of this configuration are the username and group lines. The username is the domain administrator account. The group binding allows only the group with the name `GRP` to access.



The `dn` used here is as an example only. On your network use your own domain name.

To configure LDAP server - CLI:

```
config user ldap
  edit "ldap_server"
    set server "192.168.201.3"
    set cnid "sAMAccountName"
    set dn "DC=example,DC=com,DC=au"
    set type regular
    set username "CN=Administrator,CN=Users,DC=example,DC=COM"
    set password *
    set group-member-check group-object
    set group-object-filter (&
      (objectcategory=group)member="CN=GRP,OU=training,DC=example,DC=COM")
  next
end
```

To configure the user group and add the LDAP server - CLI:

```
config user group
  edit "ldap_grp"
    set member "ldap"
    config match
      edit 1
        set server-name "ldap_server"
        set group-name "CN=GRP,OU=training,DC=example,DC=COM"
      next
    end
  next
end
```


Configuring the admin account

The wildcard part of this example is only available in the CLI for admin configuration. When enabled, this allows all LDAP group members to login to the FortiGate unit without the need to create a separate admin account for each user. In effect the members of that group will each be able to login as “test”.

To configure the admin account - CLI:

```
config system admin
  edit "test"
    set remote-auth enable
    set accprofile "super_admin"
    set wildcard enable
    set remote-group "ldap_grp"
  next
end
```

For troubleshooting, test that the admin account is operational, and see [Troubleshooting LDAP on page 171](#).

Example of LDAP to allow dial-in through member-attribute - CLI

In this example, users defined in MicroSoft Windows Active Directory (AD) are allowed to setup a VPN connection simply based on an attribute that is set to TRUE, instead of based on being part of a specific group.

In AD, the “Allow Dial-In” property is activated in the user properties, and this sets the `msNPAllowDialin` attribute to “TRUE”.

This same procedure can be used for other member attributes, as your system requires.

Configuring LDAP member-attribute settings

To accomplish this with a FortiGate unit, the member attribute must be set. Setting member attributes can only be accomplished through the CLI using the `member-attr` keyword - the option is not available through the web-based manager.

Before configuring the FortiGate unit, the AD server must be configured and have the `msNPAllowDialin` attribute set to “TRUE” for the users in question. If not, those users will not be able to properly authenticate.

The dn used here is as an example only. On your network use your own domain name.

To configure user LDAP member-attribute settings - CLI:

```
config user ldap
  edit "ldap_server"
    set server "192.168.201.3"
    set cnid "sAMAccountName"
    set dn "DC=fortinet,DC=com,DC=au"
    set type regular
    set username "fortigate@example.com"
    set password *****
    set member-attr "msNPAllowDialin"
  next
end
```

Configuring LDAP group settings

A user group that will use LDAP must be configured. This example adds the member `ldap` to the group which is the LDAP server name that was configured earlier.

To configure LDAP group settings - CLI:

```
config user group
  edit "ldap_grp"
    set member "ldap"
    config match
    edit 1
      set server-name "ldap"
      set group-name "TRUE"
    next
  end
end
```

Once these settings are in place, users can authenticate.

Troubleshooting LDAP

The examples in this section use the values from the previous example.

LDAP user test

A quick way to see if the LDAP configuration is correct is to run a `diagnose CLI` command with LDAP user information. The following command tests with a user called `netAdmin` and a password of `fortinet`. If the configuration is correct the test will be successful.

```
FGT# diag test authserver ldap ldap_server netAdmin fortinet
```

'ldap_server' is not a valid ldap server name — an LDAP server by that name has not been configured on the FortiGate unit, check your spelling.

authenticate 'netAdmin' against 'ldap_server' failed! — the user `netAdmin` does not exist on `ldap_server`, check your spelling of both the user and sever and ensure the user has been configured on the FortiGate unit.

LDAP authentication debugging

For a more in-depth test, you can use a `diag debug` command. The sample output from a shows more information about the authentication process that may prove useful if there are any problems.

Ensure the “Allow Dial-in” attribute is still set to “TRUE” and run the following CLI command. `fnbamd` is the Fortinet non-blocking authentication daemon.

```
FGT# diag debug enable
FGT# diag debug reset
FGT# diag debug application fnbamd -1
FGT# diag debug enable
```

The output will look similar to:

```
get_member_of_groups-Get the memberOf groups.
get_member_of_groups- attr='msNPAllowDialin', found 1 values
```

```
get_member_of_groups-val[0]='TRUE'  
fnbamd_ldap_get_result-Auth accepted  
fnbamd_ldap_get_result-Going to DONE state res=0  
fnbamd_auth_poll_ldap-Result for ldap svr 192.168.201.3 is SUCCESS  
fnbamd_auth_poll_ldap-Passed group matching
```

If the “Allow Dial-in” attribute is not set but it is expected, the last line of the above output will instead be:

```
fnbamd_auth_poll_ldap-Failed group matching
```

TACACS+ servers

When users connect to their corporate network remotely, they do so through a remote access server. As remote access technology has evolved, the need for security when accessing networks has become increasingly important. This need can be filled using a Terminal Access Controller Access-Control System (TACACS+) server.

TACACS+ is a remote authentication protocol that provides access control for routers, network access servers, and other networked computing devices via one or more centralized servers. TACACS+ allows a client to accept a username and password and send a query to a TACACS+ authentication server. The server host determines whether to accept or deny the request and sends a response back that allows or denies the user access to the network.

TACACS+ offers fully encrypted packet bodies, and supports both IP and AppleTalk protocols. TACACS+ uses TCP port 49, which is seen as more reliable than RADIUS’s UDP protocol.

There are several different authentication protocols that TACACS+ can use during the authentication process:

Authentication protocols

Protocol	Definition
ASCII	Machine-independent technique that uses representations of English characters. Requires user to type a username and password that are sent in clear text (unencrypted) and matched with an entry in the user database stored in ASCII format.
PAP	Password Authentication Protocol (PAP) Used to authenticate PPP connections. Transmits passwords and other user information in clear text.
CHAP	Challenge-Handshake Authentication Protocol (CHAP) Provides the same functionality as PAP, but is more secure as it does not send the password and other user information over the network to the security server.
MS-CHAP	MicroSoft Challenge-Handshake Authentication Protocol v1 (MSCHAP) Microsoft-specific version of CHAP.
default	The default protocol configuration, Auto, uses PAP, MS-CHAP, and CHAP, in that order.

Configuring a TACACS+ server on the FortiGate unit

A maximum of 10 remote TACACS+ servers can be configured for authentication.

One or more servers must be configured on FortiGate before remote users can be configured. To configure remote users, see [Local and remote users on page 182](#).



The TACACS+ page in the web-based manager (**User & Device > TACACS+ Servers**) is not available until a TACACS+ server has been configured in the CLI. For more information see the CLI Reference.

To configure the FortiGate unit for TACACS+ authentication - web-based manager:

1. Go to **User & Device > TACACS+ Servers** and select **Create New**.
2. Enter the following information, and select **OK**.

Name	Enter the name of the TACACS+ server.
Server Name/IP	Enter the server domain name or IP address of the TACACS+ server.
Server Key	Enter the key to access the TACACS+ server.
Authentication Type	Select the authentication type to use for the TACACS+ server. Auto tries PAP, MSCHAP, and CHAP (in that order).

To configure the FortiGate unit for TACACS+ authentication - CLI:

```
config user tacacs+
  edit "TACACS-SERVER"
    set server [IP_ADDRESS]
    set key [PASSWORD]
    set authen-type ascii
  next
end
config user group
  edit "TACACS-GROUP"
    set group-type firewall
    set member "TACACS-SERVER"
  next
end
config system admin
  edit TACACS-USER
    set remote-auth enable
    set accprofile "super_admin"
    set vdom "root"
    set wildcard enable
    set remote-group "TACACS-GROUP"
  next
end
```

IPv6 TACACS+ server IP address

IPv6 address support is available for TACACS+ servers.

Syntax

```
config user tacacs+
  edit <name>
    set server <ipv6 address>
    set source-ipv6 <ipv6 address>
  next
```

```
end
```

POP3 servers

FortiOS can authenticate users who have accounts on POP3 or POP3s email servers. POP3 authentication can be configured only in the CLI.

To configure the FortiGate unit for POP3 authentication:

```
config user pop3
  edit pop3_server1
    set server pop3.fortinet.com
    set secure starttls
    set port 110
  end
```

To configure a POP3 user group:

```
config user group
  edit pop3_grp
    set member pop3_server1
  end
```

A user group can list up to six POP3 servers as members.

SSO servers

Novell and Microsoft Windows networks provide user authentication based on directory services: eDirectory for Novell, Active Directory for Windows. Users can log on at any computer in the domain and have access to resources as defined in their user account. The Fortinet Single Sign On (FSSO) agent enables FortiGate units to authenticate these network users for security policy or VPN access without asking them again for their username and password.

When a user logs in to the Windows or Novell domain, the FSSO agent sends to the FortiGate unit the user's IP address and the names of the user groups to which the user belongs. The FortiGate unit uses this information to maintain a copy of the domain controller user group database. Because the domain controller authenticates users, the FortiGate unit does not perform authentication. It recognizes group members by their IP address.

In the FortiOS FSSO configuration, you specify the server where the FSSO Collector agent is installed. The Collector agent retrieves the names of the Novell or Active Directory user groups from the domain controllers on the domains, and then the FortiGate unit gets them from the Collector agent. You cannot use these groups directly. You must define FSSO type user groups on your FortiGate unit and then add the Novell or Active Directory user groups to them. The FSSO user groups that you created are used in security policies and VPN configurations to provide access to different services and resources.

FortiAuthenticator servers can replace the Collector agent when FSSO is using polling mode. The benefits of this is that FortiAuthenticator is a stand-alone server that has the necessary FSSO software pre-installed. For more information, see the [FortiAuthenticator Administration Guide](#).

SSO agent configuration settings

The following are SSO configuration settings in **Security Fabric > Fabric Connectors**.

SSO server List

Lists all the collector agents' lists that you have configured (along with other Security Fabric connectors). On this page, you can create, edit or delete FSSO agents. There are different types of FSSO agents, each with its own settings.



You can create a redundant configuration on your unit if you install a collector agent on two or more domain controllers. If the current (or first) collector agent fails, the Fortinet unit switches to the next one in its list of up to five collector agents.

Create New	Gives you the option to create a new agent. When you select Create New , you are automatically redirected to the New Fabric Connector page. Select an option from under SSO/Identity .
Edit	<p>Modifies the settings for the selected SSO server.</p> <p>To remove multiple entries from the list, for each servers you want removed, select the check box and then select Delete.</p> <p>To remove all agents from the list, on the FSSO Agent page, select the check box at the top of the check box column and then select Delete.</p>
Delete	Removes an agent from the list on the page.

Settings for Poll Active Directory Server

Server IP/Name	The IP address of the domain controller (DC).
User	The user ID used to access the domain controller.
Password	Enter the password for the account used to access the DC.
LDAP Server	Select the check box and select an LDAP server to access the Directory Service.
Enable Polling	Enable to allow the FortiGate unit to poll this DC.
Users/Groups	A list of user and user group names retrieved from the DC.

Settings when Type is RADIUS Single Sign On Agent

Name	Enter a name for the SSO server.
Use RADIUS Shared Secret	Enable and specify the SSO server secret.
Send RADIUS Responses	Enable to send RADIUS responses.

Settings for Fortinet Single Sign On Agent

Name	Enter a name for the SSO server.
Primary FSSO Agent	Enter the IP address or name of the Directory Service server where this SSO agent is installed, along with the password. The maximum number of characters is 63.
FSSO Agent	Optionally, add and configured up to four additional FSSO agents, up to a maximum of five.
Collector Agent AD access mode	Select one of the following options: <ul style="list-style-type: none"> • Standard: Enable and view A list of user and user group names retrieved from the server. • Advanced: Enable and select an LDAP server to access the Directory Service.

RSA ACE (SecurID) servers

SecurID is a two-factor system that uses one-time password (OTP) authentication. It is produced by the company RSA. This system includes portable tokens carried by users, an RSA ACE/Server, and an Agent Host. In our configuration, the FortiGate unit is the Agent Host.

Components

When using SecurID, users carry a small device or “token” that generates and displays a pseudo-random password. According to RSA, each SecurID authenticator token has a unique 64-bit symmetric key that is combined with a powerful algorithm to generate a new code every 60 seconds. The token is time-synchronized with the SecurID RSA ACE/Server.

The RSA ACE/Server is the management component of the SecurID system. It stores and validates the information about the SecurID tokens allowed on your network. Alternately the server could be an RSA SecurID 130 Appliance.

The Agent Host is the server on your network, in this case it is the FortiGate unit, that intercepts user logon attempts. The Agent Host gathers the user ID and password entered from their SecurID token, and sends that information to the RSA ACE/Server to be validated. If valid, a reply comes back indicating it is a valid logon and the FortiGate unit allows the user access to the network resources specified in the associated security policy.

Configuring the SecurID system

To use SecurID with a FortiGate unit, you need:

- to configure the RSA server and the RADIUS server to work with each other (see RSA server documentation)
- [to configure the RSA SecurID 130 Appliance](#)
- or
- [to configure the FortiGate unit as an Agent Host on the RSA ACE/Server](#)
- [to configure the FortiGate unit to use the RADIUS server](#)
- [to create a SecurID user group](#)
- [to configure a security policy with SecurID authentication](#)

The following instructions are based on RSA ACE/Server version 5.1, or RSA SecurID 130 Appliance, and assume that you have successfully completed all the external RSA and RADIUS server configuration steps listed above.

For this example, the RSA server is on the internal network, with an IP address of 192.128.100.100. The FortiGate unit internal interface address is 192.168.100.3, RADIUS shared secret is fortinet123, RADIUS server is at IP address 192.168.100.102.

To configure the RSA SecurID 130 appliance

1. Go to the IMS Console for SecurID and logon.
2. Go to **RADIUS > RADIUS Clients**, and select **Add New**.
3. Enter the following information to configure your FortiGate as a SecurID Client, and select Save.

RADIUS Client Basics	
Client Name	FortiGate
Associated RSA Agent	FortiGate
RADIUS Client Settings	
IP Address	192.168.100.3 The IP address of the FortiGate unit internal interface.
Make / Model	Select Standard Radius
Shared Secret	fortinet123 The RADIUS shared secret.
Accounting	Leave unselected
Client Status	Leave unselected

To configure the FortiGate unit as an Agent Host on the RSA ACE/Server

1. On the RSA ACE/Server computer, go to **Start > Programs > RSA ACE/Server**, and then **Database Administration - Host Mode**.
2. On the **Agent Host** menu, select **Add Agent Host**.
3. Enter and save the following information.

Name	FortiGate
Network Address	192.168.100.3 The IP address of the FortiGate unit.
Secondary Nodes	Optionally enter other IP addresses that resolve to the FortiGate unit.

If needed, refer to the RSA ACE/Server documentation for more information.

To configure the FortiGate unit to use the RADIUS server

1. Go to **User & Device > RADIUS Servers** and select **Create New**.
2. Enter the following information, and select **OK**.

Name	RSA
Primary Server IP/Name	192.168.100.102 Optionally select Test to ensure the IP address is correct and the FortiGate can contact the RADIUS server.
Primary Server Secret	fortinet123
Authentication Scheme	Select Use Default Authentication Scheme .

To create a SecurID user group

1. Go to **User & Device > User Groups**, and select **Create New**.
2. Enter the following information.

Name	RSA_group
Type	Firewall

3. In **Remote Groups**, select **Add**, then select the RSA server.
4. Select **OK**.

To create a SecurID user:

1. Go to **User & Device > User Definition**, and select **Create New**.
2. Use the wizard to enter the following information, and then select **Create**.

User Type	Remote RADIUS User
User Name	wloman
RADIUS Server	RSA
Contact Info	(optional) Enter Email or SMS information
User Group	RSA_group

To test this configuration, on your FortiGate unit use the CLI command:

```
diagnose test authserver radius RSA auto wloman 111111111
```

The series of 1s is the one time password that your RSA SecurID token generates and you enter.

Using the SecurID user group for authentication

You can use the SecurID user group in several FortiOS features that authenticate by user group including

- [Security policy](#)
- [IPsec VPN XAuth](#)
- [PPTP VPN](#)
- [SSL VPN](#)

The following sections assume the SecurID user group is called `securIDgrp` and has already been configured. Unless otherwise states, default values are used.

Security policy

To use SecurID in a security policy, you must include the SecurID user group in a security policy. This procedure will create a security policy that allows HTTP, FTP, and POP3 traffic from the `internal` interface to `wan1`. If these interfaces are not available on your FortiGate unit, substitute other similar interfaces.

To configure a security policy with SecurID authentication

1. Go to **Policy & Objects > IPv4 Policy**.
2. Select **Create New**.
3. Enter:

Incoming Interface	internal
Source Address	all
Source User(s)	securIDgrp
Outgoing Interface	wan1
Destination Address	all
Schedule	always
Services	HTTP, FTP, POP3
Action	ACCEPT
NAT	On
Shared Shaper	On, if you want to either limit traffic or guarantee minimum bandwidth for traffic that uses the SecurID security policy. Use the default shaper guarantee-100kbps .
Log Allowed Traffic	On, if you want to generate usage reports on traffic authenticated with this policy.

4. Select **OK**.
The SecurID security policy is configured.

For more detail on configuring security policies, see the FortiOS Handbook FortiGate Fundamentals guide.

IPsec VPN XAuth

Extended Authentication (XAuth) increases security by requiring user authentication in addition to the pre-shared key.

When creating an IPsec VPN using the wizard, under **VPN > IPsec Wizard**, select the SecurID **User Group** on the Authentication page. Members of the SecurID group are required to enter their SecureID code to authenticate.

For more on XAuth, see [Configuring XAuth authentication on page 230](#)

PPTP VPN

PPTP VPN is configured in the CLI. In the PPTP configuration (`config vpn pptp`), set `usrgrp` to the SecurID user group.

SSL VPN

You need to map the SecurID user group to the portal that will serve SecurID users and include the SecurID user group in the **Source User(s)** field in the security policy.

To map the SecurID group to an SSL VPN portal:

1. Go to **VPN > SSL-VPN Settings**.
2. In **Authentication/Portal Mapping**, select **Create New**.
3. Enter

Users/Groups	securIDgrp
Portal	Choose the portal.

4. Select **OK**.

Users and user groups

FortiGate authentication controls system access by user group. By assigning individual users to the appropriate user groups you can control each user's access to network resources. The members of user groups are user accounts, of which there are several types. Local users and peer users are defined on the FortiGate unit. User accounts can also be defined on remote authentication servers.

This section describes how to configure local users and peer users and then how to configure user groups. For information about configuration of authentication servers see [Authentication servers on page 155](#).

This section contains the following topics:

- [Users](#)
- [User groups](#)

Users

A user is a user account consisting of username, password, and in some cases other information, configured on the FortiGate unit or on an external authentication server. Users can access resources that require authentication only if they are members of an allowed user group. There are several different types of user accounts with slightly different methods of authentication:

User type	Authentication
Local user	The username and password must match a user account stored on the FortiGate unit. Authentication by FortiGate security policy.
Remote user	The username must match a user account stored on the FortiGate unit and the username and password must match a user account stored on the remote authentication server. FortiOS supports LDAP, RADIUS, and TACACS+ servers.
Authentication server user	A FortiGate user group can include user accounts or groups that exist on a remote authentication server.
FSSO user	With Fortinet Single Sign On (FSSO), users on a Microsoft Windows or Novell network can use their network authentication to access resources through the FortiGate unit. Access is controlled through FSSO user groups which contain Windows or Novell user groups as their members.
PKI or Peer user	A Public Key Infrastructure (PKI) or peer user is a digital certificate holder who authenticates using a client certificate. No password is required, unless two-factor authentication is enabled.
IM Users	IM users are not authenticated. The FortiGate unit can allow or block each IM user name from accessing the IM protocols. A global policy for each IM protocol governs access to these protocols by unknown users.
Guest Users	Guest user accounts are temporary. The account expires after a selected period of time.

This section includes:

- [Local and remote users](#)
- [PKI or peer users](#)
- [Two-factor authentication](#)
- [FortiToken](#)
- [Monitoring users](#)

Local and remote users

Local and remote users are defined on the FortiGate unit in **User & Device > User Definition**.

Create New	Creates a new user account. When you select Create New , you are automatically redirected to the User Creation Wizard.
Edit User	Modifies a user's account settings. When you select Edit , you are automatically redirected to the Edit User page.
Delete	<p>Removes a user from the list. Removing the user name removes the authentication configured for the user.</p> <p>The Delete icon is not available if the user belongs to a user group.</p> <p>To remove multiple local user accounts from within the list, on the User page, in each of the rows of user accounts you want removed, select the check box and then select Delete.</p> <p>To remove all local user accounts from the list, on the User page, select the check box in the check box column and then select Delete.</p>
User Name	The user name. For a remote user, this username must be identical to the username on the authentication server.
Type	Local indicates a local user authenticated on the FortiGate unit. For remote users, the type of authentication server is shown: LDAP, RADIUS, or TACACS+.
Two-factor Authentication	Indicates whether two-factor authentication is configured for the user.
Ref.	<p>Displays the number of times this object is referenced by other objects. Select the number to open the Object Usage window and view the list of referring objects. The list is grouped into expandable categories, such as Firewall Policy. Numbers of objects are shown in parentheses.</p> <p>To view more information about the referring object, use the icons:</p> <ul style="list-style-type: none"> • View the list page for these objects – available for object categories. Goes to the page where the object is listed. For example, if the category is User Groups, opens User Groups list. • Edit this object – opens the object for editing. • View the details for this object – displays current settings for the object.

To create a local or remote user account - web-based manager:

1. Go to **User & Device > User Definition** and select **Create New**.
2. On the **Choose User Type** page select:

Local User	Select to authenticate this user using a password stored on the FortiGate unit.
Remote RADIUS User Remote TACACS+ User Remote LDAP User	To authenticate this user using a password stored on an authentication server, select the type of server and then select the server from the list. You can select only a server that has already been added to the FortiGate unit configuration.

3. Select **Next** and provide user authentication information.
For a local user, enter the **User Name** and **Password**.
For a remote user, enter the **User Name** and the server name.
4. Select **Next** and enter **Contact Information**.
If email or SMS is used for two-factor authentication, provide the email address or SMS cell number at which the user will receive token password codes. If a custom SMS service is used, it must already be configured. See [FortiToken on page 188](#).
5. Select **Next**, then on the **Provide Extra Info** page enter

Two-factor Authentication	Select to enable two-factor authentication. Then select the Token (FortiToken or FortiToken Mobile) for this user account. See Associating FortiTokens with accounts on page 192 .
User Group	Select the user groups to which this user belongs.

6. Select **Create**.

To create a local user - CLI example:

Locally authenticated user

```
config user local
  edit user1
    set type password
    set passwd ljt_pj2gpepfdw
  end
```

To create a remote user - CLI example:

```
config user local
  edit user2
    set type ldap
    set ldap_server ourLDAPsrv
  end
```

For a RADIUS or TACACS+ user, set `type` to `radius` or `tacacs+`, respectively.

To create a user with FortiToken Mobile two-factor authentication - CLI example:

```
config user local
  edit user5
    set type password
    set passwd ljt_pj2gpepfdw
    set two_factor fortitoken
    set fortitoken 182937197
  end
```

Remote users are configured for FortiToken two-factor authentication similarly.

To create a user with SMS two-factor authentication using FortiGuard messaging service - CLI example:

```
config user local
  edit user6
    set type password
    set passwd 3ww_pjt68dw
    set two_factor sms
    set sms-server fortiguard
    set sms-phone 1365984521
  end
```

Removing users

Best practices dictate that when a user account is no longer in use, it should be deleted. Removing local and remote users from FortiOS involve the same steps.

If the user account is referenced by any configuration objects, those references must be removed before the user can be deleted. See [Removing references to users on page 184](#).

To remove a user from the FortiOS configuration - web-based manager:

1. Go to **User & Device > User Definition**.
2. Select the check box of the user that you want to remove.
3. Select **Delete**.
4. Select **OK**.

To remove a user from the FortiOS configuration - CLI example:

```
config user local
  delete user4444
end
```

Removing references to users

You cannot remove a user that belongs to a user group. Remove the user from the user group first, and then delete the user.

To remove references to a user - web-based manager

1. Go to **User & Device > User Definition**.
2. If the number in the far right column for the selected user contains any number other than zero, select it.

3. A more detailed list of object references to this user is displayed. Use its information to find and remove these references to allow you to delete this user.

PKI or peer users

A PKI, or peer user, is a digital certificate holder. A PKI user account on the FortiGate unit contains the information required to determine which CA certificate to use to validate the user's certificate. Peer users can be included in firewall user groups or peer certificate groups used in IPsec VPNs. For more on certificates, see [Certificates overview on page 242](#).

To define a peer user you need:

- a peer username
- the text from the subject field of the user's certificate, or the name of the CA certificate used to validate the user's certificate

Creating a peer user

The peer user can be configured only in the CLI.

To create a peer user for PKI authentication - CLI example:

```
config user peer
  edit peer1
    set subject peer1@mail.example.com
    set ca CA_Cert_1
  end
```

There are other configuration settings that can be added or modified for PKI authentication. For example, you can configure the use of an LDAP server to check access rights for client certificates. For information about the detailed PKI configuration settings, see the FortiGate CLI Reference.

Two-factor authentication

The standard logon requires a username and password. This is one factor authentication—your password is one piece of information you need to know to gain access to the system.

Two factor authentication adds the requirement for another piece of information for your logon. Generally the two factors are something you know (password) and something you have (certificate, token, etc.). This makes it harder for a hacker to steal your logon information. For example if you have a FortiToken device, the hacker would need to both use it and know your password to gain entry to your account.

Two-factor authentication is available on both user and admin accounts. But before you enable two-factor authentication on an administrator account, you need to ensure you have a second administrator account configured to guarantee administrator access to the FortiGate unit if you are unable to authenticate on the main admin account for some reason.



Two-factor authentication does not work with explicit proxies.

FortiToken extension to comply with PCI 3.2

With `multi-factor-authentication` enabled as `mandatory` (see syntax below), all authentication will collect both username/password and OTP as a second factor before presenting an authentication result. The system will log for each factor.

If a user is not configured with two-factor authentication, any OTP or an empty OTP would make the second factor authentication pass.

FortiOS processes the user and password first and then always collects the second factor (if configured) without any indication of the first factor failing or succeeding. FortiOS accepts the second factor even if the first failed (unknown to the user) and returns a login attempt pass or fail, with no indication of which factor failed.

Syntax

```
config system global
  set multi-factor-authentication {optional | mandatory}
end
```

The methods of two-factor authentication include:

- [Certificate](#)
- [Email](#)
- [SMS](#)
- [FortiToken](#)

Certificate

You can increase security by requiring both certificate and password authentication for PKI users. Certificates are installed on the user's computer. Requiring a password also protects against unauthorized use of that computer.

Optionally peer users can enter the code from their FortiToken instead of the certificate.

To create a peer user with two-factor authentication - CLI example

```
config user peer
  edit peer1
    set subject E=peer1@mail.example.com
    set ca CA_Cert_1
    set two-factor enable
    set passwd fdktguefheygfe
  end
```

For more information on certificates, see [Certificates overview on page 242](#).

Email

Two-factor email authentication sends a randomly generated six digit numeric code to the specified email address. Enter that code when prompted at logon. This token code is valid for 60 seconds. If you enter this code after that time, it will not be accepted.

A benefit is that you do not require mobile service to authenticate. However, a potential issue is if your email server does not deliver the email before the 60 second life of the token expires.

The code will be generated and emailed at the time of logon, so you must have email access at that time to be able to receive the code.

To configure an email provider - web-based manager:

1. Go to **System > Advanced** and enable **Use Custom Email Server** under **Email Service**.
2. Enter **SMTP Server** and **Default Reply To** address.
3. If applicable, enable **Authentication** and enter the **SMTP User** and **Password** to use.
4. Select a **Security Mode**, options are: **None**, **SMTPS** or **STARTTLS**.
5. Enter the **Port** number, the default is 25.
6. Select **Apply**.

To configure an email provider - CLI:

```
config system email-server
    set server <server_domain-name>
    set reply-to <Recipient_email_address>
end
```

To enable email two-factor authentication - web-based manager:

1. To modify an administrator account, go to **System > Administrators**. To modify a user account go to **User & Device > User Definition**.
2. Edit the user account.
3. Enable and enter the user's **Email Address**.
4. Select **Enable Two-factor Authentication**.
5. Select **Email based two-factor authentication**.
6. Select **OK**.

If **Email based two-factor authentication** option doesn't appear after selecting **Enable Two-factor Authentication**, you need to enable it via the CLI as follows.



To enable email two-factor authentication - CLI:

```
config user local
    edit <user_name>
        set email-to <user_email>
        set two-factor email
    end
```

SMS

SMS two-factor authentication sends the token code in an SMS text message to the mobile device indicated when this user attempts to logon. This token code is valid for 60 seconds. If you enter this code after that time, it will not be accepted. Enter this code when prompted at logon to be authenticated.

SMS two-factor authentication has the benefit that you do not require email service before logging on. A potential issue is if the mobile service provider does not send the SMS text message before the 60 second life of the token expires.

FortiGuard Messaging Service include four SMS Messages at no cost. If you need more, you should acquire a license through support.fortinet.com or via customer service.

If you do not use the FortiGuard Messaging Service, you need to configure an SMS service.

To configure an SMS service - CLI:

```
config system sms-server
  edit <provider_name>
    set mail-server <server_domain-name>
  next
end
```

To configure SMS two-factor authentication - web-based manager:

1. To modify an:
 - administrator account, go to **System > Administrators**, or
 - user account go to **User & Device > User Definition**.
2. Edit the user account.
3. Select **SMS** and enter the **Country Dial Code** and **Phone Number**.
4. Select **Enable Two-factor Authentication**, and select the correct **Token**.
5. Select **OK**.

If you have problems receiving the token codes via SMS messaging, contact your mobile provider to ensure you are using the correct phone number format to receive text messages and that your current mobile plan allows text messages.

FortiToken

FortiToken is a disconnected one-time password (OTP) generator. It is a small physical device with a button that when pressed displays a six digit authentication code. This code is entered with a user's username and password as two-factor authentication. The code displayed changes every 60 seconds, and when not in use the LCD screen is blanked to extend the battery life.

There is also a mobile phone application, FortiToken Mobile, that performs much the same function.

FortiTokens have a small hole in one end. This is intended for a lanyard to be inserted so the device can be worn around the neck, or easily stored with other electronic devices. Do not put the FortiToken on a key ring as the metal ring and other metal objects can damage it. The FortiToken is an electronic device like a cell phone and must be treated with similar care.

Any time information about the FortiToken is transmitted, it is encrypted. When the FortiGate unit receives the code that matches the serial number for a particular FortiToken, it is delivered and stored encrypted. This is in keeping with the Fortinet's commitment to keeping your network highly secured.

FortiTokens can be added to user accounts that are local, IPsec VPN, SSL VPN, and even Administrators. See [Associating FortiTokens with accounts on page 192](#).

A FortiToken can be associated with only one account on one FortiGate unit.

If a user loses their FortiToken, it can be locked out using the FortiGate so it will not be used to falsely access the network. Later if found, that FortiToken can be unlocked on the FortiGate to allow access once again. See [FortiToken maintenance on page 194](#).

There are three tasks to complete before FortiTokens can be used to authenticate accounts:

1. [Adding FortiTokens to the FortiGate](#)
2. [Activating a FortiToken on the FortiGate](#)
3. [Associating FortiTokens with accounts](#)

In addition, this section includes the following:

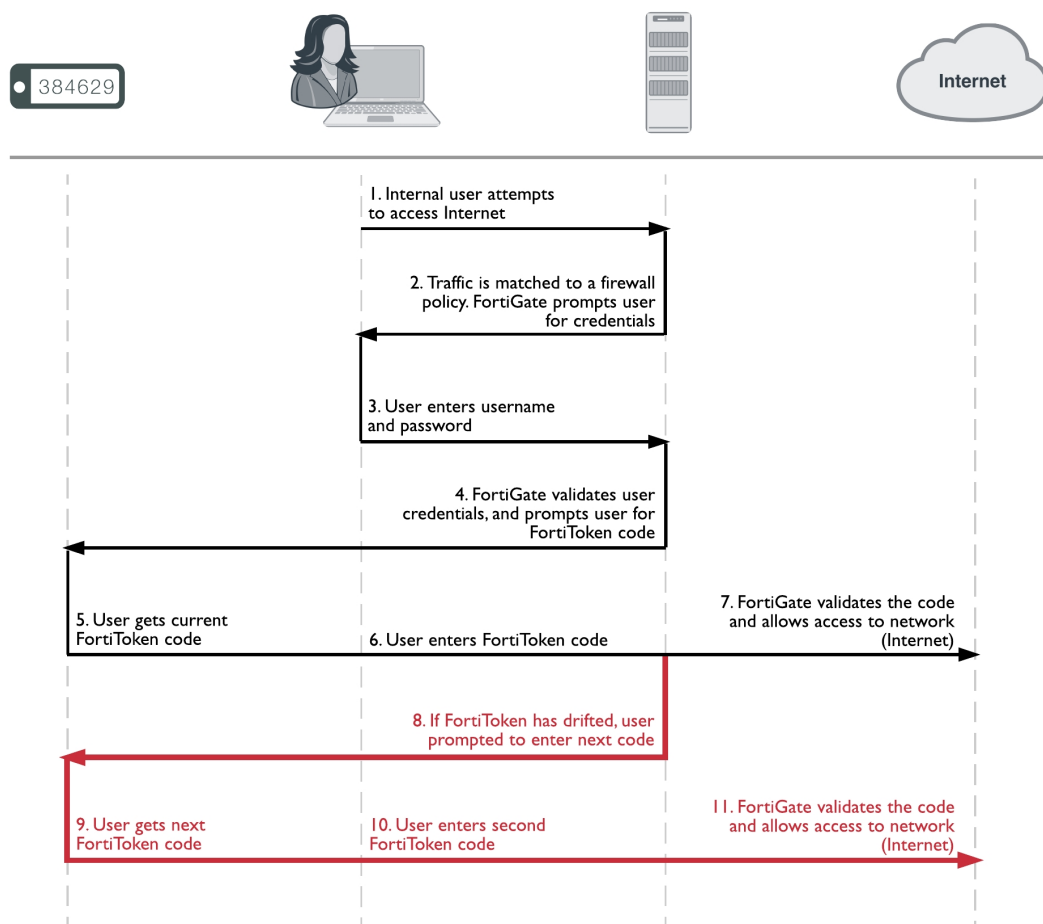
- [FortiToken maintenance](#)
- [FortiToken Mobile Push](#)

The FortiToken authentication process

The steps during FortiToken two-factor authentication are as follows.

1. User attempts to access a network resource.
2. FortiGate unit matches the traffic to an authentication security policy, and FortiGate unit prompts the user for username and password.
3. User enters their username and password.
4. FortiGate unit verifies their information, and if valid prompts the user for the FortiToken code.
5. User gets the current code from their FortiToken device.
6. User enters current code at the prompt.
7. FortiGate unit verifies the FortiToken code, and if valid allows access to the network resources such as the Internet.
The following steps are needed only if the time on the FortiToken has drifted and needs to be re-synchronized with the time on the FortiGate unit.
8. If time on FortiToken has drifted, FortiGate unit will prompt user to enter a second code to confirm.
9. User gets the next code from their FortiToken device
10. User enters the second code at the prompt.
11. FortiGate unit uses both codes to update its clock to match the FortiToken and then proceeds as in step "[Users and user groups](#)" on page 181.

The FortiToken authentication process is illustrated below:



When configured the FortiGate unit accepts the username and password, authenticates them either locally or remotely, and prompts the user for the FortiToken code. The FortiGate then authenticates the FortiToken code. When FortiToken authentication is enabled, the prompt field for entering the FortiToken code is automatically added to the authentication screens.

Even when an Administrator is logging in through a serial or Telnet connection and their account is linked to a FortiToken, that Administrator will be prompted for the token's code at each login.



If you have attempted to add invalid FortiToken serial numbers, there will be no error message. The serial numbers will simply not be added to the list.

Adding FortiTokens to the FortiGate

Before one or more FortiTokens can be used to authenticate logons, they must be added to the FortiGate. The import feature is used to enter many FortiToken serial numbers at one time. The serial number file must be a text file with one FortiToken serial number per line.



Both FortiToken Mobile and physical FortiTokens store their encryption seeds on the cloud, therefore you will only be able to register them to a single FortiGate or FortiAuthenticator.

Because FortiToken-200CD seed files are stored on the CD, these tokens can be registered on multiple FortiGates and/or FortiAuthenticators, but **not** simultaneously.

To manually add a FortiToken to the FortiGate - web-based manager:

1. Go to **User & Device > FortiTokens**.
2. Select **Create New**.
3. In **Type**, select **Hard Token** or **Mobile Token**.
4. Enter one or more FortiToken serial numbers (hard token) or activation codes (mobile token).
5. Select **OK**.



For mobile token, you receive the activation code in the license certificate once you purchase a license. FortiOS include a license for two mobile token at no cost.

To import multiple FortiTokens to the FortiGate - web-based manager:

1. Go to **User & Device > FortiTokens**.
2. Select **Create New**.
3. In **Type**, select **Hard Token**.
4. Select **Import**.
5. Select **Serial Number File** or **Seed File**, depending on which file you have.
6. Browse to the local file location on your local computer.
7. Select **OK**.
The file is imported.
8. Select **OK**.

To import FortiTokens to the FortiGate from external sources - CLI:

FortiToken seed files (both physical and mobile versions) can be imported from either FTP or TFTP servers, or a USB drive, allowing seed files to be imported from an external source more easily:

```
execute fortitoken import ftp <file name> <ip>[:ftp port] <Enter> <user> <password>
execute fortitoken import tftp <file name> <ip>
execute fortitoken import usb <file name>
```



To import seed files for FortiToken Mobile, replace `fortitoken` with `fortitoken-mobile`.

To add two FortiTokens to the FortiGate - CLI:

```
config user fortitoken
  edit <serial_number>
```

```
next
edit <serial_number2>
next
end
```

Activating a FortiToken on the FortiGate

Once one or more FortiTokens have been added to the FortiGate unit, they must be activated before being available to be associated with accounts. The process of activation involves the FortiGate querying FortiGuard servers about the validity of each FortiToken. The serial number and information is encrypted before it is sent for added security.



A FortiGate unit requires a connection to FortiGuard servers to activate a FortiToken.

To activate a FortiToken on the FortiGate unit - web-based manager:

1. Go to **User & Device > FortiTokens**.
 2. Select one or more FortiTokens with a status of Available.
 3. Right-click the FortiToken entry and select **Activate**.
 4. Select **Refresh**.
- The status of selected FortiTokens will change to Activated.

The selected FortiTokens are now available for use with user and admin accounts.

To activate a FortiToken on the FortiGate unit - CLI:

```
config user fortitoken
edit <token_serial_num>
set status activate
next
end
```

Associating FortiTokens with accounts

The final step before using the FortiTokens to authenticate logons is associating a FortiToken with an account. The accounts can be local user or administrator accounts.

To add a FortiToken to a local user account - web-based manager:

1. Ensure that your FortiToken serial number has been added to the FortiGate successfully, and its status is Available.
2. Go to **User & Device > User Definition**, and edit the user account.
3. Enter the user's **Email Address**.
4. Enable **Two-factor Authentication**.
5. Select the user's FortiToken serial number from the **Token** list.
6. Select **OK**.



For mobile token, click on **Send Activation Code** to be sent to the email address configured previously. The user will use this code to activate his mobile token. An **Email Service** has to be set under **System > Advanced** in order to send the activation code.

To add a FortiToken to a local user account - CLI:

```
config user local
  edit <username>
    set type password
    set passwd "myPassword"
    set two-factor fortitoken
    set fortitoken <serial_number>
    set email-to "username@example.com"
    set status enable
  next
end
```

To add a FortiToken to an administrator account - web-based manager:

1. Ensure that your FortiToken serial number has been added to the FortiGate successfully, and its status is Available.
2. Go to **System > Administrators**, and edit the admin account.
This account is assumed to be configured except for two-factor authentication.
3. Enter admin's **Email Address**.
4. Enable **Two-factor Authentication**.
5. Select the user's FortiToken serial number from the **Token** list.
6. Select **OK**.



For mobile token, click on **Send Activation Code** to be sent to the email address configured previously. The admin will use this code to activate his mobile token. An **Email Service** has to be set under **System > Advanced** in order to send the activation code.

To add a FortiToken to an administrator account - CLI:

```
config system admin
  edit <username>
    set password "myPassword"
    set two-factor fortitoken
    set fortitoken <serial_number>
    set email-to "username@example.com"
  next
end
```

The `fortitoken` keyword will not be visible until `fortitoken` is selected for the `two-factor` option.



Before a new FortiToken can be used, it may need to be synchronized due to clock drift.

FortiToken maintenance

Once FortiTokens are entered into the FortiGate unit, there are only two tasks to maintain them — changing the status,

To change the status of a FortiToken between activated and locked - CLI:

```
config user fortitoken
  edit <token_serial_num>
    set status lock
  next
end
```

Any user attempting to login using this FortiToken will not be able to authenticate.

To list the drift on all FortiTokens configured on this FortiGate unit - CLI:

```
# diag fortitoken info
FORTITOKEN DRIFT STATUS
FTK2000BHV1KRZCC 0 token already activated, and seed won't be returned
FTK2001C5YCRRVEE 0 token already activated, and seed won't be returned
FTKMOB4B94972FBA 0 provisioned
FTKMOB4BA4BE9B84 0 new
Total activated token: 0
Total global activated token: 0
Token server status: reachable
```

This command lists the serial number and drift for each FortiToken configured on this FortiGate unit. This command is useful to check if it is necessary to synchronize the FortiGate and any particular FortiTokens.

FortiToken Mobile Push

A command under `config system ftm-push` allows you to configure the FortiToken Mobile Push services server IP address and port number. The Push service is provided by Apple (APNS) and Google (GCM) for iPhone and Android smartphones respectively. This will help to avoid tokens becoming locked after an already enabled two-factor authentication user has been disabled.

CLI syntax

```
config system ftm-push
  set server-ip <ip-address>
  set server-port [1-65535] Default is 4433.
end
```

Note that the `server-ip` is the public IP address of the FortiGate interface that the FTM will call back to; it is the IP address used by the FortiGate for incoming FTM calls.

If an SSL VPN user authenticates with their token, then logs out and attempts to reauthenticate again within a minute, a new message will display showing "Please wait x seconds to login again." This replaces a previous error/permission denied message.

The "x" value will depend on the calculation of how much time is left in the current time step.

CLI syntax

```
config system interface
edit <name>
set allowaccess ftm
next
end
```



FortiGate supports when the FortiAuthenticator initiates FTM Push notifications, for when users are attempting to authenticate through a VPN and/or RADIUS (with FortiAuthenticator as the RADIUS server).

Monitoring users

To monitor user activity in the web-based manager, go to **Monitor > Firewall User Monitor**. The list of users who are logged on is displayed with some information about them such as their user group, security policy ID, how long they have been logged on, their IP address, traffic volume, and their authentication method as one of FSSO, NTLM, or firewall (FW-auth).

From this screen you can de-authenticate all users who are logged on. The de-authenticate button is at the top left of this screen.

To see information about banned users go to **Monitor > Quarantine Monitor**. Displayed information about users who have been banned includes what application triggered the ban (Application Protocol), the reason for the ban (Cause or rule), Created, and when the ban expires.

Filtering the list of users

When there are many users logged on, it can be difficult to locate a specific user or multiple users to analyze. Applying filters to the list allows you to organize the user list to meet your needs, or only display some the users that meet your current requirements.

Select settings bottom at the top right of the screen to adjust columns that are displayed for users, including what order they are displayed in. This can be very helpful in locating information you are looking for.

Each column heading has a grey filter icon. Click on the filter icon to configure a filter for the data displayed in that column. Each column has similar options including a field to enter the filtering information, a check box to select the negative of the text in the field, and the options to add more fields, apply the filter, clear all filters, or cancel without saving. To enter multiple terms in the field, separate each of them with a comma. To filter entries that contain a specific prefix, use an * (asterisk).

For example, to create a filter to display only users with an IP address of 10.11.101.x who authenticated using one of security policies five through eight, and who belong to the user group Accounting.

1. Go to **Monitor > Quarantine Monitor**.
2. Enter 10.11.101.* and select **Apply**.
3. Select the filter icon beside **Policy ID**.
4. Enter 5-8 and select **Apply**.
5. Select the filter icon beside **User Group**.
6. Enter Accounting and select **Apply**.

**Login credentials for guest users shown in clear text on GUI and voucher**

In FortiOS 5.6.4, login credentials for guest users is displayed/printed in clear text on the GUI and in the voucher. It is also sent in clear text by SMS and email.

User groups

A user group is a list of user identities. An identity can be:

- a local user account (username/password stored on the FortiGate unit)
- a remote user account (password stored on a RADIUS, LDAP, or TACACS+ server)
- a PKI user account with digital client authentication certificate stored on the FortiGate unit
- a RADIUS, LDAP, or TACACS+ server, optionally specifying particular user groups on that server
- a user group defined on an FSSO server.

Security policies and some types of VPN configurations allow access to specified user groups only. This restricted access enforces Role Based Access Control (RBAC) to your organization's network and its resources. Users must be in a group and that group must be part of the security policy.

In most cases, the FortiGate unit authenticates users by requesting their username and password. The FortiGate unit checks local user accounts first. If a match is not found, the FortiGate unit checks the RADIUS, LDAP, or TACACS+ servers that belong to the user group. Authentication succeeds when a matching username and password are found. If the user belongs to multiple groups on a server, those groups will be matched as well.



FortiOS does not allow username overlaps between RADIUS, LDAP, or TACACS+ servers.

There are four types of FortiGate user groups: Firewall, FSSO, Guest, and RADIUS single sign-on (RSSO) user groups.

Firewall user groups

Firewall user groups are used locally as part of authentication. When a security policy allows access only to specified user groups, users must authenticate. If the user authenticates successfully and is a member of one of the permitted groups, the session is allowed to proceed.

This section includes:

- [SSL VPN access](#)
- [IPsec VPN access](#)
- [Configuring a firewall user group](#)
- [Multiple group enforcement support](#)
- [User group timeouts](#)

SSL VPN access

SSL VPN settings include a list of the firewall user groups that can access the SSL VPN and the SSL VPN portal that each group will use. When the user connects to the FortiGate unit via HTTPS on the SSL VPN port (default 10443), the FortiGate unit requests a username and password.

SSL VPN access also requires a security policy where the destination is the SSL interface. For more information, see the [FortiOS Handbook SSL VPN](#) guide.

IPsec VPN access

A firewall user group can provide access for dialup users of an IPsec VPN. In this case, the IPsec VPN phase 1 configuration uses the **Accept peer ID in dialup group** peer option. The user's VPN client is configured with the username as peer ID and the password as pre-shared key. The user can connect successfully to the IPsec VPN only if the username is a member of the allowed user group and the password matches the one stored on the FortiGate unit.



A user group cannot be used as a dialup group if any member of the group is authenticated using an external authentication server.

For more information, see the [FortiOS Handbook IPsec VPN](#) guide.

Configuring a firewall user group

A user group can contain:

- local users, whether authenticated by the FortiGate unit or an authentication server
- PKI users
- authentication servers, optionally specifying particular user groups on the server

To create a Firewall user group - web-based manager:

1. Go to **User & Device > User Groups** and select **Create New**.
2. Enter a name for the user group.
3. In **Type**, select **Firewall**.
4. Add user names to the **Members** list.
5. Add authentication servers to the **Remote groups** list.

By default all user accounts on the authentication server are members of this FortiGate user group. To include only specific user groups from the authentication server, deselect **Any** and enter the group name in the appropriate format for the type of server. For example, an LDAP server requires LDAP format, such as: `cn=users,dn=office,dn=example,dn=com`

Remote servers must already be configured in **User & Device**.

6. Select **OK**.

To create a firewall user group - CLI example:

In this example, the members of `accounting_group` are `User1` and all of the members of `rad_accounting_group` on myRADIUS external RADIUS server.

```
config user group
  edit accounting_group
    set group-type firewall
    set member User1 myRADIUS
  config match
    edit 0
      set server-name myRADIUS
      set group-name rad_accounting_group
```

```
end
end
```



Matching user group names from an external authentication server might not work if the list of group memberships for the user is longer than 8000 bytes. Group names beyond this limit are ignored.

`server_name` is the name of the RADIUS, LDAP, or TACACS+ server, but it must be a member of this group first and must also be a configured remote server on the FortiGate unit.

`group_name` is the name of the group on the RADIUS, LDAP, or TACACS+ server such as “engineering” or “cn=users,dc=test,dc=com”.

Before using group matching with TACACS+, you must first enable authentication. For example if you have a configured TACACS+ server called myTACS, use the following CLI commands.

```
config user tacacs+
  edit myTACS
    set authorization enable
  next
end
```

For more information about user group CLI commands, see the [Fortinet CLI Guide](#).

Multiple group enforcement support

Previously, when a user belonged to multiple user groups, this user could only access the group services that were within one group. With multiple group enforcement, a user can access the services within the groups that the user is part of.

For example, `userA` belongs to `user_group1`, `user_group2`, `user_group3`, and `user_group4`; previously `userA` could only access services within one of those four groups, typically the group that matches the first security policy. This can be annoying if HTTP access is in `user_group1`, FTP access is in `user_group2`, and email access is in `user_group3`. Now `userA` can access services within `user_group1`, `user_group2`, `user_group3`, and `user_group4`.

This feature is available only in the CLI and is enabled by default. It applies to RADIUS, LDAP, and TACACS+ servers. The new command for this feature is `auth-multi-group` found in `config user settings` and checks all groups a user belongs to for authentication.

User group timeouts

User groups can have timeout values per group in addition to FortiGate-wide timeouts. There are essentially three different types of timeouts that are configurable for user authentication on the FortiGate unit — idle timeout, hard timeout, and session timeout. These are in addition to any external timeouts such as those associated with RADIUS servers.

If VDOMs are enabled, the global level user setting `authtimeout` is the default all VDOMs inherit. If VDOMs are not enabled, user settings `authtimeout` is the default. The default timeout value is used when the `authtimeout` keyword for a user group is set to zero.

Each type of timeout will be demonstrated using the existing user group `example_group`. Timeout units are minutes. A value of zero indicates the global timeout is used.

Membership in multiple groups

When a user belongs to multiple groups in RADIUS groups, the group auth-timeout values are ignored. Instead the global timeout value is used. The default value is 5 minutes, but it can be set from 1 to 43200 minutes (30 days).

```
config user setting
    set auth-timeout-type idle-timeout
    set auth-timeout 300
end
```

Idle timeout

The default type of timeout is idle timeout. When a user initiates a session, it starts a timer. As long as data is transferred in this session, the timer continually resets. If data flow stops, the timer is allowed to advance until it reaches its limit. At that time the user has been idle for too long, and the user is forced to re-authenticate before traffic is allowed to continue in that session.

To configure user group authentication idle timeout - CLI:

```
config user settings
    set auth-timeout-type idle-timeout
end
config user group
    edit example_group
        set authtimeout 5 //range is 0-43200 minutes (0 = use global authtimeout value)
    next
end
```

Hard timeout

Where the idle timeout is reset with traffic, the hard timeout is absolute. From the time the first session a user establishes starts, the hard timeout counter starts. When the timeout is reached, all the sessions for that user must be re-authenticated. This timeout is not affected by any event.

To configure user group authentication hard timeout - CLI:

```
config user settings
    set auth-timeout-type hard-timeout
end
config user group
    edit example_group
        set authtimeout 43200 //range is 0-43200 minutes (0 = use global authtimeout value)
    next
end
```

Session timeout

The session timeout works much like the hard timeout in that its an absolute timer that can not be affected by events. However, when the timeout is reached existing sessions may continue but new sessions are not allowed until re-authentication takes place.

To configure a user group authentication new session hard timeout - CLI:

```
config user setting
    set auth-timeout-type new-session
end

config user group
    edit example_group
        set authtimeout 30 //range is 0-43200 minutes (0 = use global authtimeout value)
    next
end
```

SSO user groups

SSO user groups are part of FSSO authentication and contain only Windows or Novell network users. No other user types are permitted as members. Information about the Windows or Novell user groups and the logon activities of their members is provided by the Fortinet Single Sign On (FSSO) which is installed on the network domain controllers.

You can specify FSSO user groups in security policies in the same way as you specify firewall user groups. FSSO user groups cannot have SSL VPN or dialup IPsec VPN access.

For information about configuring FSSO user groups, see [Creating FSSO user groups on page 313](#). For complete information about installing and configuring FSSO, see [Agent-based FSSO on page 276](#).

Configuring peer user groups

Peer user groups can only be configured using the CLI. Peers are digital certificate holders defined using the `config user peer` command. The peer groups you define here are used in dialup IPsec VPN configurations that accept RSA certificate authentication from members of a peer certificate group.

To create a peer group - CLI

```
config user peergrp
    edit vpn_peergrp1
        set member pki_user1 pki_user2 pki_user3
    end
```

Viewing, editing, and deleting user groups

To view the list of FortiGate user groups, go to **User & Device > User Groups**.

Editing a user group

When editing a user group in the CLI you must set the type of group this will be — either a firewall group, a Fortinet Single Sign-On Service group (FSSO), a Radius based Single Sign-On Service group (RSSO), or a guest group. Once the type of group is set, and members are added you cannot change the group type without removing the members.

In the web-based manager, if you change the type of the group any members will be removed automatically.

To edit a user group - web-based manager

1. Go to **User & Device > User Groups**.
2. Select the user group that you want to edit.
3. Select the **Edit** button.
4. Modify the user group as needed.
5. Select **OK**.

To edit a user group - CLI

This example adds user3 to Group1. Note that you must re-specify the full list of users:

```
config user group
  edit Group1
    set group-type firewall
    set member user2 user4 user3
  end
```

Deleting a user group

Before you delete a user group, you must ensure there are no objects referring to, it such as security policies. If there are, you must remove those references before you are able to delete the user group.

To remove a user group - web-based manager

1. Go to **User & Device > User Groups**.
2. Select the user group that you want to remove.
3. Select the **Delete** button.
4. Select **OK**.

To remove a user group - CLI

```
config user group
  delete Group2
end
```

SSL renegotiation in firewall authentication

The `auth-ssl-allow-renegotiation` option is available under `config user setting` to allow/forbid SSL renegotiation in firewall authentication. The default value is `disable`, where a session would be terminated by `authd` once renegotiation is detected and this login would be recorded as a failure. Other behavior follows regular authentication settings.

To enable SSL renegotiation - CLI

```
config user setting
  set auth-ssl-allow-renegotiation enable
end
```


Managing guest access

Visitors to your premises might need user accounts on your network for the duration of their stay. If you are hosting a large event such as a conference, you might need to create many such temporary accounts. The FortiOS Guest Management feature is designed for this purpose.

A guest user account User ID can be the user's email address, a randomly generated string, or an ID that the administrator assigns. Similarly, the password can be administrator-assigned or randomly generated.

You can create many guest accounts at once using randomly-generated User IDs and passwords. This reduces administrator workload for large events.

User's view of guest access

1. The user receives an email, SMS message, or printout from a FortiOS administrator listing a User ID and password.
2. The user logs onto the network with the provided credentials.
3. After the expiry time, the credentials are no longer valid.

Administrator's view of guest access

1. Create one or more guest user groups.
All members of the group have the same characteristics: type of User ID, type of password, information fields used, type and time of expiry.
2. Create guest accounts using Guest Management.
3. Use captive portal authentication and select the appropriate guest group.

Configuring guest user access

To set up guest user access, you need to create at least one guest user group and add guest user accounts. Optionally, you can create a guest management administrator whose only function is the creation of guest accounts in specific guest user groups. Otherwise, any administrator can do guest management.

Creating guest management administrators

To create a guest management administrator

1. Go to **System > Administrators** and create a regular administrator account.
For detailed information see the [System Administration](#) chapter.
2. Select **Restrict to Provision Guest Accounts**.
3. In **Guest Groups**, add the guest groups that this administrator manages.

Creating guest user groups

The guest group configuration determines the fields that are provided when you create a guest user account.

To create a guest user group:

1. Go to **User & Device > User Groups** and select **Create New**.
2. Enter the following information:

Name	Enter a name for the group.
Type	Guest
Enable Batch Account Creation	<p>Create multiple accounts automatically. When this is enabled:</p> <ul style="list-style-type: none"> • User ID and Password are set to Auto-Generate. • The user accounts have only User ID, Password, and Expiration fields. Only the Expiration field is editable. If the expiry time is a duration, such as “8 hours”, this is the time after first login. • You can print the account information. Users do not receive email or SMS notification. <p>See To create multiple guest user accounts automatically on page 204.</p>
User ID	<p>Select one of:</p> <ul style="list-style-type: none"> • Email — User’s email address • Specify — Administrator assigns user ID • Auto-Generate — FortiGate unit creates a random user ID
Password	<p>Select one of:</p> <ul style="list-style-type: none"> • Specify — Administrator assigns user ID • Auto-Generate — FortiGate unit creates a random password • Disable — no password
Expire Type	<p>Choose one of:</p> <ul style="list-style-type: none"> • Immediately — expiry time is counted from creation of account • After first login — expiry time is counted from user’s first login
Default Expire Time	Set the expire time. The administrator can change this for individual users.
Enable Name	If enabled, user must provide a name.
Enable Sponsor	If enabled, user form has Sponsor field. Select Required if required.
Enable Company	If enabled, user form has Company field. Select Required if required.
Enable Email	If enabled, user is notified by email.
Enable SMS	If enabled, user is notified by SMS. Select whether FortiGuard Messaging Service or a another SMS provider is used.

Creating guest user accounts

Guest user accounts are not the same as local user accounts created in **User & Device > User Definition**. Guest accounts are not permanent; they expire after a defined time period. You create guest accounts in **User & Device > Guest Management**.

To create a guest user account

1. Go to **User & Device > Guest Management**.
2. In **Guest Groups**, select the guest group to manage.
3. Select **Create New** and fill in the fields in the **New User** form.

Fields marked Optional can be left blank. The guest group configuration determines the fields that are available.

4. Select **OK**.

To create multiple guest user accounts automatically

1. Go to **User & Device > Guest Management**.
2. In **Guest Groups**, select the guest group to manage.
The guest group must have the **Enable Batch Guest Account Creation** option enabled.
3. Select **Create New > Multiple Users**.
Use the down-pointing caret to the right of **Create New**.
4. Enter **Number of Accounts**.
5. Optionally, change the **Expiration**.
6. Select **OK**.

Guest management account List

Go to **User & Device > Guest Management** to create, view, edit or delete guest user accounts.

Create New	Creates a new guest user account.
Edit	Edit the selected guest user account.
Delete	Delete the selected guest user account.
Purge	Remove all expired accounts from the list.
Send	Send the user account information to a printer or to the guest. Depending on the group settings and user information, the information can be sent to the user by email or SMS.
Refresh	Update the list.
Guest Groups	Select the guest group to list. New accounts are added to this group.
User ID	The user ID. Depending on the guest group settings, this can be the user's email address, an ID that the administrator specified, or a randomly-generated ID.
Expires	Indicates a duration such as "3 hours". A duration on its own is relative to the present time. Or, the duration is listed as "after first login."

Guest access in a retail environment

Some retail businesses such as coffee shops provide free WiFi Internet access for their customers. For this type of application, the FortiOS guest management feature is not required; the WiFi access point is open and customers do not need logon credentials. However, the business might want to contact its customers later with promotional offers to encourage further patronage. Using an Email Collection portal, it is possible to collect customer email addresses for this purpose. The security policy grants network access only to users who provide a valid email address.

The first time a customer's device attempts to use the WiFi connection, FortiOS requests an email address, which it validates. The customer's subsequent connections go directly to the Internet without interruption.

Creating an email harvesting portal

The customer's first contact with your network will be with a captive portal which presents a web page requesting an email address. When FortiOS has validated the email address, the customer's device MAC address is added to the Collected Emails device group.

To create the email collection portal:

1. Go to **WiFi & Switch Controller > SSID** and edit your SSID.
2. Set **Security Mode** to **Captive Portal**.
3. Set **Portal Type** to **Email Collection**.
4. Optionally, in **Customize Portal Messages** select **Email Collection**.

You can change the portal content and appearance. See [Customizing captive portal pages on page 235](#).

To create the email collection portal - CLI:

In this example the `freewifi` WiFi interface is modified to present an email collection captive portal.

```
config wireless-controller vap
  edit freewifi
    set security captive-portal
    set portal-type email-collect
  end
```

Creating the security policy

You need configure a security policy that allows traffic to flow from the WiFi SSID to the Internet interface but only for members of the Collected Emails device group. This policy must be listed first. Unknown devices are not members of the Collected Emails device group, so they do not match the policy.

To create the security policy:

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information:

Incoming Interface	freewifi
Source Address	all
Source Device Type	Collected Emails
Outgoing Interface	wan1
Destination Address	all
Schedule	always
Service	ALL

Action	ACCEPT
NAT	On

3. Select **OK**.

To create the authentication rule - CLI:

```
config firewall policy
edit 3
    set srcintf "freewifi"
    set dstintf "wan1"
    set srcaddr "all"
    set action accept
    set devices collected-emails
    set nat enable
    set schedule "always"
    set service "ALL"
end
```

Checking for harvested emails

In the web-based manager, go to **User & Device > Device Inventory**. In the CLI you can use the `diagnose user device list` command. For example,

```
FGT-100D # diagnose user device list
hosts
vd 0 d8:d1:cb:ab:61:0f gen 35 req 30 redir 1 last 43634s 7-11_2-int
ip 10.0.2.101 ip6 fe80::dad1:cbff:feab:610f
type 2 'iPhone' src http c 1 gen 29
os 'iPhone' version 'iOS 6.0.1' src http id 358 c 1
email 'yo@yourdomain.com'
vd 0 74:e1:b6:dd:69:f9 gen 36 req 20 redir 0 last 39369s 7-11_2-int
ip 10.0.2.100 ip6 fe80::76e1:b6ff:fedd:69f9
type 1 'iPad' src http c 1 gen 5
os 'iPad' version 'iOS 6.0' src http id 293 c 1
host 'Joes's-iPad' src dhcp
email 'you@fortinet.com'
```

Fall-through authentication policies

User authentication policies have an implicit fall-through feature that intentionally causes policy matching to fall through to a policy lower on the list that can also match the traffic. In other words the first user policy that is matched in the policy list, based on standard policy criteria, isn't the only policy that can be matched.

Fall-through is intended to match users in different user groups with different policies. For example, consider an organization with two user groups where one group requires a web filtering profile, while the other requires virus scanning. In this example, you would edit two basic Internet access policies: policy 1 assigning **User Group A** with a **Web Filtering** profile, and policy 2 assigning **User Group B** with an **AntiVirus** profile. Both policies are also assigned to the same internal subnet, named **subnet1**.

In this configuration, all users from **subnet1** will see an authentication prompt. If the user is found in **User Group A**, the traffic is accepted by policy 1 and is filtered by the **Web Filtering** profile. If the user is found in **User Group B**, the traffic is accepted by policy 2 and is virus scanned.

The fall-through feature is required for users to be matched with policy 2. Without the fall-through feature, traffic would never be matched with policy 2.

Configuring authenticated access

When you have configured authentication servers, users, and user groups, you are ready to configure security policies and certain types of VPNs to require user authentication.

This section describes:

- [Authentication timeout](#)
- [Password policy](#)
- [Authentication protocols](#)
- [Authentication in captive portals](#)
- [Authentication in security policies](#)
- [Authentication replacement messages](#)
- [VPN authentication](#)

Authentication timeout

An important feature of the security provided by authentication is that it is temporary—a user must re-authenticate after logging out. Also if a user is logged on and authenticated for an extended period of time, it is a good policy to have them re-authenticate at set periods. This ensures a user's session is cannot be spoofed and used maliciously for extended periods of time — re-authentication will cut any spoof attempts short. Shorter timeout values are more secure.

Security authentication timeout

You set the security user authentication timeout to control how long an authenticated connection can be idle before the user must authenticate again. The maximum timeout is 4320 minutes (72 hours).

To set the security authentication timeout - web-based manager:

1. Go to **User & Device > Authentication Settings**.
2. Enter the **Authentication Timeout** value in minutes.
The default authentication timeout is 5 minutes.
3. Select **Apply**.

SSL VPN authentication timeout

You set the SSL VPN user authentication timeout (**Idle Timeout**) to control how long an authenticated connection can be idle before the user must authenticate again. The maximum timeout is 259 200 seconds. The default timeout is 300 seconds.

To set the SSL VPN authentication timeout - web-based manager:

1. Go to **VPN > SSL-VPN Settings**.
2. Enable **Idle Logout** and enter the **Inactive For** value in seconds.
3. Select **Apply**.

Password policy

Password authentication is effective only if the password is sufficiently strong and is changed periodically. By default, the FortiGate unit requires only that passwords be at least eight characters in length, but up to 128 characters is permitted. You can set a password policy to enforce higher standards for both length and complexity of passwords. Password policies can apply to administrator passwords or IPsec VPN pre-shared keys.

To set a password policy in the web-based manager, go to **System > Settings**. In the CLI, use the `config system password-policy` command.

Users usually create passwords composed of alphabetic characters and perhaps some numbers. Password policy can require the inclusion of uppercase letters, lowercase letters, numerals or punctuation characters.

Configuring password minimum requirement policy

Best practices dictate that passwords include:

- one or more uppercase characters
- one or more lower case characters
- one or more of the numerals
- one or more special characters.

The minimum number of each of these types of characters can be set in both the web-based manager and the CLI.

The following procedures show how to force administrator passwords to contain at least two uppercase, four lower case, two digits, and one special character. Leave the minimum length at the default of eight characters.

To change administrator password minimum requirements - web-based manager:

1. Go to **System > Settings**.
2. Select **Enable Password Policy**.
3. Select **Must Contain at Least**.
4. Enter the following information:

Upper Case Letters	2
Lower Case Letters	4
Numbers	2
Special Characters	1

5. Under **Apply Password Policy to**, select **Administrator Password**.
6. Select **Apply**.

To change administrator password minimum requirements - CLI:

```
config system password-policy
  set status enable
  set apply-to admin-password
  set min-upper-case-letter 2
  set min-lower-case-letter 4
```



```
set min-number 2
set min-non-alphanumeric 1
set change-4-characters enable
end
```

The `change-4-characters` option forces new passwords to change a minimum of four characters in the old password. Changing fewer characters results in the new password being rejected. This option is only available in the CLI.

To configure a guest administrator password policy - CLI:

As of FortiOS 5.4, a password policy can also be created for guest administrators. The following command shows all possible commands, which are also available under `config system password-policy`.

```
config system password-policy
  set status {enable | disable} Enable/disable password policy.
  set apply-to {guest-admin-password} Guest admin to which this password policy applies.
  set minimum-length <8-128> Minimum password length.
  set min-lower-case-letter <0-128> Min. lowercase characters in password.
  set min-upper-case-letter <0-128> Min. uppercase characters in password.
  set min-non-alphanumeric <0-128> Min. non-alphanumeric characters in password.
  set min-number <0-128> Min. numeric characters in password.
  set change-4-characters {enable | disable} Enable/disable changing at least 4 characters for new password.
  set expire-status {enable | disable} Enable/disable password expiration.
  set expire-day <1-999> Number of days before password expires.
  set reuse-password {enable | disable} Enable/disable reuse of password.
end
```

Password best practices

In addition to length and complexity, there are security factors that cannot be enforced in a policy. Guidelines issued to users will encourage proper password habits.

Best practices dictate that password expiration also be enabled. This forces passwords to be changed on a regular basis. You can set the interval in days. The more sensitive the information this account has access to, the shorter the password expiration interval should be. For example 180 days for guest accounts, 90 days for users, and 60 days for administrators.

Avoid:

- real words found in any language dictionary
- numeric sequences, such as “12345”
- sequences of adjacent keyboard characters, such as “qwerty”
- adding numbers on the end of a word, such as “hello39”
- adding characters to the end of the old password, such as “hello39” to “hello3900”
- repeated characters
- personal information, such as your name, birthday, or telephone number.



In the case where the FortiGate is handling multiple keytabs in Kerberos authentication, use different passwords when generating each keytab.

Maximum login attempts and blackout period

When you login and fail to enter the correct password you could be a valid user, or a hacker attempting to gain access. For this reason, best practices dictate to limit the number of failed attempts to login before a blackout period where you cannot login.

To set a maximum of five failed authentication attempts before the blackout, using the following CLI command:

```
config user setting
    set auth-invalid-max 5
end
```

To set the length of the blackout period to five minutes, or 300 seconds, once the maximum number of failed login attempts has been reached, use the following CLI command:

```
config user setting
    set auth-blackout-time 300
end
```

Authentication protocols

When user authentication is enabled on a security policy, the authentication challenge is normally issued for any of the four protocols, HTTP, HTTPS, FTP, and Telnet, which are dependent on the connection protocol. By making selections in the Protocol Support list, the user controls which protocols support the authentication challenge. The user must connect with a supported protocol first, so that they can subsequently connect with other protocols.

For example, if you have selected HTTP, FTP, or Telnet, a username and password-based authentication occurs. The FortiGate unit then prompts network users to input their security username and password. If you have selected HTTPS, certificate-based authentication (HTTPS, or HTTP redirected to HTTPS only) occurs.



FTP and Telnet authentication replacement messages cannot be customized. For HTTP and HTTPS replacement messages see [Authentication replacement messages on page 213](#).

For certificate-based authentication, you must install customized certificates on the FortiGate unit and on the browsers of network users. If you do not install certificates on the network user's web browser, the network users may see an SSL certificate warning message and have to manually accept the default FortiGate certificate. The network user's web browser may deem the default certificate as invalid.

When you use certificate authentication, if you do not specify any certificate when you create the security policy, the global settings are used. If you specify a certificate, the per-policy setting will overwrite the global setting. For more information about the use of certification authentication see [Certificate-based authentication on page 241](#).

To set the authentication protocols

1. Go to **User & Device > Authentication Settings**.
2. In **Protocol Support**, select the required authentication protocols.
3. If using HTTPS protocol support, in **Certificate**, select a Local certificate from the drop-down list.
4. Select **Apply**.

Authentication in captive portals

Network interfaces, including WiFi interfaces, can perform authentication at the interface level using a captive portal — an HTML form that requests the user's name and password. A captive portal is useful where all users connecting to the network interface must authenticate. Optionally, on a WiFi interface, the captive portal can be combined with a terms of service disclaimer to which the user must agree before gaining access. For more information, see [Captive portals on page 233](#).

Once successfully authenticated, the user's session passes to the firewall.

Authentication in security policies

Security policies control traffic between FortiGate interfaces, both physical interfaces and VLAN subinterfaces. The firewall tries to match the session's user or group identity, device type, destination, or other attribute to a security policy. When a match is found, the user connects to the requested destination. If no security policy matches, the user is denied access.

A user who has not already been authenticated by a captive portal, FSSO, or RSSO can match only policies where no user or user group is specified. If no such policy exists, the firewall requests authentication. If the user can authenticate and the session can be matched to a policy, the user connects to the requested destination, otherwise, the user is denied access.

This section includes:

- [Enabling authentication protocols](#)
- [Authentication replacement messages](#)
- [Access to the Internet](#)
- [Configuring authentication security policies](#)
- [Identity-based policy](#)
- [NTLM authentication](#)
- [Certificate authentication](#)
- [Restricting number of concurrent user logins](#)

Enabling authentication protocols

Users can authenticate using FTP, HTTP, HTTPS, and Telnet. However, these protocols must be enabled first.

Another authentication option is to redirect any attempts to authenticate using HTTP to a more secure channel that uses HTTPS. This forces users to a more secure connection before entering their user credentials.

To enable support for authentication protocols - web-based manager:

1. Go to **User & Device > Authentication Settings**.
2. Select one or more of HTTP, HTTPS, FTP, Telnet, or Redirect HTTP Challenge to a Secure Channel (HTTPS). Only selected protocols will be available for use in authentication.
3. Select the **Certificate** to use, for example `Fortinet_Factory`.
4. Select **Apply**.

To enable support for authentication protocols - CLI:

```
config user setting
```

```

set auth-type ftp http https telnet
set auth-cert Fortinet_Factory
end

```



As of FortiOS 5.4, the `Fortinet_Factory` certificate has been re-signed with an expiration date of 2038. It is used instead of `Fortinet_Factory2`, which has been removed.

Authentication replacement messages

A replacement message is the body of a web page containing a message about a blocked website message, a file too large message, a disclaimer, or even a login page for authenticating. The user is presented with this message instead of the blocked content.

Authentication replacement messages are the prompts a user sees during the security authentication process such as login page, disclaimer page, and login success or failure pages. These are different from most replacement messages because they are interactive requiring a user to enter information, instead of simply informing the user of some event as other replacement messages do.

Replacement messages have a system-wide default configuration, a per-VDOM configuration, and disclaimers can be customized for multiple security policies within a VDOM.

These replacement messages are used for authentication using HTTP and HTTPS. Authentication replacement messages are HTML messages. You cannot customize the security authentication messages for FTP and Telnet.

The authentication login page and the authentication disclaimer include replacement tags and controls not found on other replacement messages.

More information about replacement messages can be found in the `config system replacemsg` section of the [FortiOS CLI Reference](#).

List of authentication replacement messages

Replacement message name (CLI name)	Description
Login challenge page (auth-challenge-page)	<p>This HTML page is displayed if security users are required to answer a question to complete authentication. The page displays the question and includes a field in which to type the answer. This feature is supported by RADIUS and uses the generic RADIUS challenge-access auth response. Usually, challenge-access responses contain a Reply-Message attribute that contains a message for the user (for example, "Please enter new PIN"). This message is displayed on the login challenge page. The user enters a response that is sent back to the RADIUS server to be verified.</p> <p>The Login challenge page is most often used with RSA RADIUS server for RSA SecurID authentication. The login challenge appears when the server needs the user to enter a new PIN. You can customize the replacement message to ask the user for a SecurID PIN.</p> <p>This page uses the <code>%%QUESTION%%</code> tag.</p>

Replacement message name (CLI name)	Description
Disclaimer page (auth-disclaimer-page-1) (auth-disclaimer-page-2) (auth-disclaimer-page-3)	<p>This page prompts user to accept the displayed disclaimer when leaving the captive portal to access Internet resources. It is displayed when the captive portal type is Authentication and Disclaimer or Disclaimer Only.</p> <p>In the CLI, the auth-disclaimer-page-2 and auth-disclaimer-page-3 pages seamlessly extend the size of the disclaimer page from 8 192 characters to 16 384 and 24 576 characters respectively. In the web-based manager this is handled automatically.</p> <p>See Disclaimer on page 216.</p>
Email token page (auth-email-token-page)	<p>The page prompting a user to enter their email token. See Email on page 1.</p>
FortiToken page (auth-fortitoken-page)	<p>The page prompting a user to enter their FortiToken code. See FortiToken on page 188.</p>
Keepalive page (auth-keepalive-page)	<p>The HTML page displayed with security authentication keepalive is enabled using the following CLI command:</p> <pre>config system globalset auth-keepalive enable end</pre> <p>Authentication keepalive keeps authenticated firewall sessions from ending when the authentication timeout ends. In the web-based manager, go to User & Device > Authentication Settings to set the Authentication Timeout.</p> <p>This page includes %%TIMEOUT%%.</p>
Login failed page (auth-login-failed-page)	<p>The Disclaimer page replacement message does not re-direct the user to a redirect URL or the security policy does not include a redirect URL. When a user selects the button on the disclaimer page to decline access through the FortiGate unit, the Declined disclaimer page is displayed.</p>
Login page (auth-login-page)	<p>The authentication HTML page displayed when users who are required to authenticate connect through the FortiGate unit using HTTP or HTTPS.</p> <p>Prompts the user for their username and password to login.</p> <p>This page includes %%USERNAMEID%% and %%PASSWORDID%% tags.</p>
Declined disclaimer page (auth-reject-page)	<p>The page displayed if a user declines the disclaimer page. See Disclaimer on page 216.</p>

Replacement message name (CLI name)	Description
SMS Token page (auth-sms-token-page)	The page prompting a user to enter their SMS token. See SMS on page 187 .
Success message (auth-success-msg)	The page displayed when a user successfully authenticates. Prompts user to attempt their connection again (as the first was interrupted for authentication).

Access to the Internet

A policy for accessing the Internet is similar to a policy for accessing a specific network, but the destination address is set to **all**. The destination interface is the one that connects to the Internet Service Provider (ISP). For general purpose Internet access, the Service is set to ALL.

Access to HTTP, HTTPS, FTP and Telnet sites may require access to a domain name service. DNS requests do not trigger authentication. You must configure a policy to permit unauthenticated access to the appropriate DNS server, and this policy must **precede** the policy for Internet access. Failure to do this will result in the lack of a DNS connection and a corresponding lack of access to the Internet.

Configuring authentication security policies

To include authentication in a security policy, the policy must specify user groups. A security policy can authenticate by certificate, FSSO, and NTLM. The two exceptions to this are RADIUS SSO and FSSO Agents. See [SSO using RADIUS accounting records on page 321](#), and [Introduction to FSSO agents on page 277](#).

Before creating a security policy, you need to configure one or more users or user groups. For more information, see [Users and user groups on page 181](#).

Creating the security policy is the same as a regular security policy except you must select the action specific to your authentication method:

Authentication methods allowed for each policy Action

Action	Authentication method	Where authentication is used
ACCEPT	FSSO Agent or a security policy that specifies an FSSO user group	Agent-based FSSO on page 276 .
	NTLM	See NTLM authentication on page 217 .
	Certificates	See Configuring certificate-based authentication on page 255 .
	RADIUS SSO	See SSO using RADIUS accounting records on page 321 .
DENY	none	none

Disclaimer

A WiFi or SSL captive portal can include a disclaimer message presented after the user authenticates. The user must agree to the terms of the disclaimer to access network resources.

Customizing authentication replacement messages

Customizing disclaimers or other authentication replacement messages involves changing the text of the disclaimer message, and possibly the overall appearance of the message.

Changing the disclaimer in **System > Replacement Messages** is not the same as selecting to customize a disclaimer used in a captive portal. The captive portal location is a customized disclaimer that inherits the default format for the disclaimer message, but then can be customized for this portal.

To customize the disclaimer for a captive portal - web-based manager:

1. Go to **Network > Interfaces**. Either select an existing interface or create a new one.
2. Under **Security Mode**, select **Captive Portal**, and enable **Customize Portal Messages**.
3. Select the **Edit** icon. You can select and edit any of the pages. Change your text or layout as needed.

Enabling security logging

There are two types of logging that relate to authentication — event logging, and security logging.

When enabled, event logging records system events such as configuration changes, and authentication. To configure event logging, go to **Log & Report > Log Settings** and enable **Event Logging**. Select the events you want to log, such as **User activity event**.

When enabled, security logging will log security profile and security policy traffic.

You must enable logging within a security policy, as well as the options that are applied to a security policy, such as security profiles features. Event logs are enabled within the Event Log page.

For more information on logging, see the FortiOS Log and Reporting guide.

For more information on specific types of log messages, see the FortiOS Log Message Reference.



You need to set the logging severity level to **Notification** when configuring a logging location to record traffic log messages.

To enable logging within an existing security policy - web-based manager:

1. Go to **Policy & Objects > IPv4 Policy**.
2. Expand to reveal the policy list of a policy.
3. Select the security policy you want to enable logging on and then select **Edit**.
4. To log all general firewall traffic, select the check box beside **Log Allowed Traffic**, and choose to enable **Security Events** or **All Sessions**.
5. Select **OK**.

Identity-based policy

An identity-based policy (IBP) performs user authentication in addition to the normal security policy duties. If the user does not authenticate, access to network resources is refused. This enforces Role Based Access Control (RBAC) to your organization's network and resources.

Identity-based policies also support Single Sign-On operation. The user groups selected in the policy are of the Fortinet Single Sign-On (FSSO) type.

User authentication can occur through any of the following supported protocols, including: HTTP, HTTPS, FTP, and Telnet. The authentication style depends on which of these protocols is included in the selected security services group and which of those enabled protocols the network user applies to trigger the authentication challenge.

For username and password-based authentication (HTTP, FTP, and Telnet) the FortiGate unit prompts network users to enter their username, password, and token code if two-factor authentication is selected for that user account. For certificate-based authentication, including HTTPS or HTTP redirected to HTTPS only, see [Certificate authentication on page 228](#).

With identity-based policies, the FortiGate unit allows traffic that matches the source and destination addresses, device types, and so on. This means specific security policies must be placed **before** more general ones to be effective.

When the identity-based policy has been configured, the option to customize authentication messages is available. This allows you to change the text, style, layout, and graphics of the replacement messages associated with this firewall policy. When enabled, customizing these messages follows the same method as changing the disclaimer. See [Disclaimer on page 216](#).

Types of authentication also available in identity-based policies are

- [NTLM authentication](#)
- [Certificate authentication](#)

NTLM authentication

NT LAN Manager (NTLM) protocol can be used as a fallback for authentication when the Active Directory (AD) domain controller is unreachable. NTLM uses the web browser to send and receive authentication information. See "NTLM" and "FSSO NTLM authentication support".

To enable NTLM

1. Edit the policy in the CLI to enable NTLM. For example, if the policy ID is 4:
2. Go to **Policy & Objects > IPv4 Policy** and note the **ID** number of your FSSO policy.
3. The policy must have an FSSO user group as **Source User(s)**. There must be at least one FSSO Collector agent configured on the FortiGate unit.

```
config firewall policy
edit 4
set ntlm enable
end
```


NTLM guest access

Guest profile access may be granted to users who fail NTLM authentication, such as visitors who have no user credentials on the network. To allow guest user access, edit the FSSO security policy in the CLI, like this:

```
config firewall policy
  edit 4
    set ntlm enable
    set ntlm-guest enable
  end
```

NTLM enabled browsers - CLI

User agent strings for NTLM enabled browsers allow the inspection of initial HTTP-User-Agent values, so that non-supported browsers are able to go straight to guest access without needlessly prompting the user for credentials that will fail. `ntlm-guest` must be enabled to use this option.

```
config firewall policy
  edit 4
    set ntlm enable
    set ntlm-guest enable
    set ntlm-enabled-browsers <user_agent_string>
  next
end
```

`<user_agent_string>` is the name of the browser that is NTLM enabled. Examples of these values include "MSIE", "Mozilla" (which includes FireFox), and "Opera".

Value strings can be up to 63 characters in length, and may not contain cross site scripting (XSS) vulnerability characters such as brackets. The FortiGate unit prevents use of these characters to prevent exploit of cross site scripting (XSS) vulnerabilities.

Kerberos authentication for explicit web and transparent web proxy users

Kerberos authentication is a method for authenticating both explicit web proxy and transparent web proxy users. It has several advantages over NTLM challenge response:

- Does not require FSSO/AD agents to be deployed across domains.
- Requires fewer round-trips than NTLM SSO, making it less latency sensitive.
- Is (probably) more scalable than challenge response.
- Uses existing Windows domain components rather than added components.
- NTLM may still be used as a fallback for non-Kerberos clients.

Enhancements to Kerberos explicit and transparent web proxy

FortiOS 5.6.x authentication is managed by schemes and rules based on protocol and source address. As such, configurable authentication settings have been introduced to enhance authentication.

CLI commands (`config authentication rule, scheme, and setting`) allow explicit proxy rules and schemes to be created to separate user authentication (e.g. authentication rules and schemes used to match conditions in order to identify users) from user authorization (proxy-based policies with users and/or user groups).

CLI syntax - config authentication rule

```

config authentication rule
  edit <name>
    set name <name>
    set status {enable|disable}
    set protocol {http|ftp|socks}
    config srcaddr <addr-name or addrgrp-name>
      edit <name>
        set name <ipv4-policy-name>
      next
    end
    config srcaddr6 <addr-name or addrgrp-name>
      edit <name>
        set name <ipv6-policy-name>
      next
    end
    set ip-based {enable|disable}
    set active-auth-method <scheme-name>
    set sso-auth-method <scheme-name>
    set transaction-based {enable|disable} - basic scheme + session-based
    set web-auth-cookie {enable|disable}
    set comments <comments>
  next
end

```

Note: As shown above, HTTP, FTP, and SOCKSv5 authentication protocols are supported for explicit proxy.

Authentication rules are used to receive user-identity, based on the values set for protocol and source address. Having said this, if a rule fails to match based on source address, there will be no other attempt to match the rule, however the next policy will be attempted. This occurs only when:

- there is an authentication rule, but no authentication method has been set (under `config authentication scheme`; see below), so user identity cannot be found.
- the user is successfully matched in the rule, but fails to match the current policy.

Once a rule is positively matched through protocol and/or source address, it must also match the authentication method specified (`active-auth-method` and `sso-auth-method`). These methods point to schemes, as defined under `config authentication scheme`.

CLI syntax - config authentication scheme

```

config authentication scheme
  edit <name>
    set name <name>
    set method {basic|digest|ntlm|form|negotiate|fsso|rsso}
    set negotiate-ntlm {enable|disable}
    set require-tfa {enable|disable}
    set fsso-guest {enable|disable}
    config user-database
      edit <name>
        set name {local|<ldap-server>|<radius-server>|<fsso-name>|<rsso-name>|<tacacs+-name>}
      next
    end
  next
end

```

```
end
```

Combining authentication rules and schemes, granular control can be exerted over users and IPs, creating an efficient process for users to successfully match a criteria before matching the policy.

Additional options can be set under `config authentication setting`.

CLI syntax - config authentication setting

```
config authentication setting
  set sso-scheme <scheme-name>
  set active-scheme <scheme-name>
  set captive-portal <host-name>
  set captive-portal-port <tcp-port>
end
```

Integration of transparent and explicit proxy HTTP policy checking

A CLI command, under `config firewall profile-protocol-options`, allows HTTP policy checking to be enable or disabled. When enabled, transparent traffic can be matched in a firewall policy and policy user authentication can occur. In addition, separate SSL inspection policies can be created:

```
config firewall profile-protocol-options
  edit <name>
    set http-policy {enable|disable}
  end
```

Internet Service Database in Explicit/Implicit proxy policies

CLI commands, under `config firewall proxy-policy`, implement the Internet Service Database (ISDB) as the webproxy matching factor, and override IP pool is also support:

```
config firewall proxy-policy
  edit <name>
    set proxy {explicit-web|transparent-web|ftp|wanopt}
    set dstintf <dst-name>
    set poolname <ip-pool-name>
  end
```

Multiple port/port range support for explicit web and explicit FTP proxy

Multiple port numbers and/or ranges can be set for explicit proxy, specifically for HTTP/HTTPS and FTP. Go to **Network > Explicit Proxy** and configure settings under **Explicit Web Proxy** and **Explicit FTP Proxy**, or under `config web-proxy explicit` in the CLI Console.

1. General configuration

1.1 Kerberos environment - Windows server setup

1. Build a Windows 2008 Platform server.
2. Enable domain configuration in windows server (dcpromo).
3. Set the domain name TEST.COM (realm name).

1.2 Create users

- *testuser* is a normal user (could be any existing domain user account).
- *testfgt* is the service name. In this case it should be the FQDN for the explicit proxy Interface, For example the hostname in the client browser proxy config.
- Recommendation: create username all in lowercase (even if against corporate standards).
 - The account only requires “domain users” membership
 - Password set to never expire
 - Set a very strong password

1.3 Add FortiGate to DNS



Add the FortiGate FQDN in to the Windows DNS domain, as well as in-addr.arpa

For Lab/Testing add the FortiGate Domain name and IP mapping in the hosts file (windows/system32/drivers/etc/hosts). e.g., `TESTFGT.TEST.COM 10.10.1.10`

1.4 Generate the Kerberos keytab

Use the *ktpass* command (found on Windows Servers and many domain workstations) to generate the Kerberos keytab.

Example:

```
ktpass -princ HTTP/<domain name of test fgt>@realm -mapuser testfgt -pass <password> -crypto all -ptype KRB5_NT_PRINCIPAL -out fgt.keytab
```



In the case where the FortiGate is handling multiple keytabs in Kerberos authentication, use different passwords when generating each keytab.



The ktpass on older Windows servers (i.e. 2003) may not support the “all” crypto option.

Example:

```
ktpass -princ HTTP/testfgt.test.com@TEST.COM -mapuser testfgt -pass 12345678 -crypto all -ptype KRB5_NT_PRINCIPAL -out fgt.keytab
```



The realm name is always presented in uppercase, and prefixed with the “@” character.

1.5 Encode base64

Use the `base64` command (available in most Linux distros) command to encode the `fgt.keytab` file. Any LF (Line Feed) need to be deleted from the file.

Example:

```
base64 fgt.keytab > fgt.txt
```



Use Notepad++ or some native Linux text editor. Windows Notepad and Wordpad are likely to introduce errors.

2. FortiGate configuration

2.1 Create LDAP server instance

```
config user ldap
  edit "ldap" <<< Required for authorization
    set server "10.10.1.1" <<< LDAP server IP, normally it should be same as KDC server
    set cnid "cn"
    set dn "dc=test,dc=com"
    set type regular
    set username "CN=admin,CN=Users,DC=test,DC=com" <<< Your domain may require STARTTLS
    set password <FOOS>
  next
end
```

2.2 Define Kerberos as an authentication service

```
config user krb-keytab
  edit "http_service"
    set principal "HTTP/testfgt.test.com@TEST.COM" <<< Same as the principal name in 1.4
    set ldap-server "ldap" <<< the defined ldap server for authorization
    set keytab
      "BQIAAABNAAIACKJFUkJFUj5DT00ABEhUVFAAGlRPT1lfRkdUXzEwMERfQS5CRVJCVRViuQ09NAAAAQA
      AAAAKABcAEJQl0MHgovwplu7XzfENJzw=" <<< base64 encoding keytab data, created in step 1.5
    next
  end
```

2.3 Create user group(s)

```
config user group <<< the group is used for kerberos authentication
  edit "testgrp"
    set member "ldap"
    config match
      edit 1
        set server-name "ldap" <<< Same as ldap-server option in krb-keytab
        set group-name "CN=Domain Users,CN=Users,DC=TEST,DC=com"
      next
    end
  next
end
```

2.4 Create firewall policy

```
config firewall proxy-policy
edit 1
set uuid 5e5dd6c4-952c-51e5-b363-120ad77c1414
set proxy explicit-web
set dstintf "port1"
set srcaddr "all"
set dstaddr "all"
set service "webproxy"
set action accept
set schedule "always"
set groups "CN=USERS LAB.PS FSSO"
next
end
```

2.5 Diagnostics

Once the keytab is imported, check that it has been properly decoded. The filename generated will be relatively random, but should be clearly visible.

```
Artoo-Deetoo (root) # fnsysctl ls -la /tmp/kt
drwxr--r-- 2 0 0 Fri Dec 2 10:06:43 2016 60 .
drwxrwxrwt 22 0 0 Tue Dec 6 14:28:29 2016 3280 ..
-rw-r--r-- 1 0 0 Fri Dec 2 10:06:43 2016 392 1.0.89.keytab
```



If there is no file present, then the file hasn't decoded. Check the file for line feeds and try again.

3. Client side walkthrough

3.1 Check Kerberos is working

Log on to the domain by using *testuser*, created in 1.2. Use the *klist* command to list ticket information. In the below example, the client has received *krbtgt*, *CIFS*, and *LDAP* tickets. As there has been no interaction with the FortiGate, there are no references to it.

```
C:\Users\glenk>klist Cached Tickets: (5)

C:\Users\glenk>klist
Cached Tickets: (5)
#0> Client: glenk @ home.local

    Server: krbtgt/HOME.LOCAL @ HOME.LOCAL
    KerbTicket Encryption Type: AES-256-CTS-HMAC-SHA1-96
    Ticket Flags 0x60a00000 -> forwardable forwarded renewable pre_authent
    Start Time: 12/6/2016 14:58:06 (local)
    End Time: 12/7/2016 0:58:04 (local)
    Renew Time: 12/13/2016 14:58:04 (local)
    Session Key Type: AES-256-CTS-HMAC-SHA1-96

#1> Client: glenk @ home.local

    Server: krbtgt/HOME.LOCAL @ HOME.LOCAL
    KerbTicket Encryption Type: AES-256-CTS-HMAC-SHA1-96
```

```

Ticket Flags 0x40e00000 -> forwardable renewable initial pre_authent
Start Time: 12/6/2016 14:58:04 (local)
End Time: 12/7/2016 0:58:04 (local)
Renew Time: 12/13/2016 14:58:04 (local)
Session Key Type: AES-256-CTS-HMAC-SHA1-96

#2> Client: glenk @ home.local

Server: cifs/EthicsGradient.home.local @ HOME.LOCAL
Kerberos Ticket Encryption Type: AES-256-CTS-HMAC-SHA1-96
Ticket Flags 0x40a40000 -> forwardable renewable pre_authent ok_as_delegate
Start Time: 12/6/2016 14:58:06 (local)
End Time: 12/7/2016 0:58:04 (local)
Renew Time: 12/13/2016 14:58:04 (local)
Session Key Type: AES-256-CTS-HMAC-SHA1-96

#3> Client: glenk @ home.local

Server: ldap/EthicsGradient.home.local @ HOME.LOCAL
Kerberos Ticket Encryption Type: AES-256-CTS-HMAC-SHA1-96
Ticket Flags 0x40a40000 -> forwardable renewable pre_authent ok_as_delegate
Start Time: 12/6/2016 14:58:06 (local)
End Time: 12/7/2016 0:58:04 (local)
Renew Time: 12/13/2016 14:58:04 (local)
Session Key Type: AES-256-CTS-HMAC-SHA1-96

#4> Client: glenk @ home.local

Server: LDAP/EthicsGradient.home.local/home.local @ HOME.LOCAL
Kerberos Ticket Encryption Type: AES-256-CTS-HMAC-SHA1-96
Ticket Flags 0x40a40000 -> forwardable renewable pre_authent ok_as_delegate
Start Time: 12/6/2016 14:58:06 (local)
End Time: 12/7/2016 0:58:04 (local)
Renew Time: 12/13/2016 14:58:04 (local)
Session Key Type: AES-256-CTS-HMAC-SHA1-96

```

3.2 Configure client

Set up web-proxy in browser through the FortiGate. This can be achieved via a PAC file or direct browser configuration.



Some Firefox documentation indicates that it is necessary to make manual advanced configuration changes to allow Kerberos authentication work. However, builds 48 (and possibly much earlier) require no additional configuration beyond setting of the proxy server.

3.3 Open a connection to the Internet

1. The client accesses the explicit proxy, but a *HTTP 407 Proxy Authentication Required* is returned.
2. As "Negotiate" is set, the client has knowledge of the KRB-TGT, it requests a ticket from the KDC with a *krb-tgs-req* message. This includes the REALM (HOME.LOCAL) in the *req-body* section, and the provided instances SNAME and service (in this case, HTTP/artoo-deetoo.home.local).
3. The KDC responds with a next KRB-TGS-REP.

This ticket is then available on the client.

In the example below, the ticket-granted-service has issued Ticket #2.

```
#2> Client: glenk @ home.local
```

```

Server: HTTP/artoo-deetoo.home.local @ HOME.LOCAL
KerbTicket Encryption Type: RSADSI RC4-HMAC (NT)
Ticket Flags 0x40a00000 -> forwardable renewable pre_authent
Start Time: 12/6/2016 14:59:45 (local)
End Time: 12/7/2016 0:58:04 (local)
Renew Time: 12/13/2016 14:58:04 (local)
Session Key Type: RSADSI RC4-HMAC (NT)

```

4. The conversation between the client and the proxy continues, as the client responds with the Kerberos ticket in the response.

The whole process takes less than a second to complete. The user should be visible as a FSSO logon in the Web UI.

Transparent web-proxy Kerberos authentication

Transparent web-proxy allows the FortiGate to process level 7 policy matching, even when the explicit web-proxy is not enabled on the client's browser. The transparent web-proxy policy is set in proxy-policy too. The policy matching rule is the same as the explicit web-proxy.

In the firewall policy level, transparent web-proxy is regarded as a special UTM. The HTTP/HTTPS traffic matches the firewall policy first, then traffic is redirected to the web-proxy daemon. If the transparent web-proxy feature is disabled, http-policy options in profile-protocol-options is used to enable transparent web-proxy feature.

IP-based

Kerberos authentication requires the captive portal to be an FQDN address that is resolved to a local IP address. However, it becomes more complicated to setup an FQDN address in a local user deployment. Therefore you can set the `captive-portal-type` to either use an FQDN or IP address.

1. Captive portal and the captive portal port must be configured in transparent web-proxy for support of Kerberos authentication:

```

config authentication setting
    set captive-portal-type {fqdn | ip}
    set captive-portal <fqdn-name> / <ip>
    set captive-portal-port "9998"
end

```

2. Authentication rule, scheme, and krb-keytab need to be configured for Kerberos authentication (note the `active-auth-method` scheme referenced in the rule):

```

config authentication scheme
    edit <kerberos-scheme>
        set method negotiate
        set negotiate-ntlm <enable>
        set fsso-guest <disable>
    next
end

config authentication rule
    edit <name>
        set status <enable>
        set protocol <http>
        set srcaddr "all"
        set ip-based <enable>

```



```

        set active-auth-method <kerberos-scheme>
    next
end

config user krb-keytab
    edit <name>
        set principal "HTTP/TESTFGT.TEST.COM@TEST.COM"
        set ldap-server "ldap"
        set keytab <base64-encoding-keytab-data>
    next
end

```

3. Configure LDAP and user group used for authorization:

```

config user ldap
    edit "ldap"
        set server "10.10.1.1"
        set cnid
        set dn
        set type <regular>
        set username "CN=admin,CN=Users,DC=test,DC=com"
        set password ENC
            aW5lIAHkPMf4D+ZCKpGMU3x8Fpq0G+7uIbAvpblbXFA5vLfGb4/oRBx+B6R/v+CMCetP84e+Gdz5zEcM
            yOd3cj0BoIhFrpYJfXhRs4lSE0HezeVXfxwTSf5VJG+F11G/G5RpaY+AE8bortC8MBe7P2/uGQocFHu4
            Ilulp5I60Jvyk6Ei3hDZMjTd8iPp5IkRJZVVjQ==
    next
end

config user group
    edit "testgrp"
        set member "ldap"
        config match
            edit "1"
                set server-name "ldap"
                set group-name "CN=Domain Users,CN=Users,DC=TEST,DC=com"
            next
        end
    next
end

```

4. Create proxy-policy, with groups as the authorizing policy-matching element:

```

config firewall proxy-policy
    edit 1
        set uuid 1bbb891a-9cd2-51e7-42ff-d1fa13cac3da
        set proxy explicit-web
        set dstintf "any"
        set srcaddr "all"
        set dstaddr "all"
        set service "webproxy"
        set action accept
        set schedule "always"
        set groups testgrp
    next
end

```

5. UTM must be enabled in the firewall policy to support the transparent web-proxy:

```

config firewall policy

```

```

edit "1"
    set name "policy1"
    set uuid 8a6ceeac-b016-51e6-2b5c-165070d5bf50
    set srcintf "mgmt1"
    set dstintf "mgmt1"
    set srcaddr "all"
    set dstaddr "all"
    set action <accept>
    set schedule "always"
    set service "ALL"
    set utm-status <enable>
    set profile-protocol-options "transparent-web-proxy"
    set ssl-ssh-profile "deep-inspection"
    set nat <enable>
next
end

config firewall profile-protocol-options
    edit "transparent-web-proxy"
        config http
            set ports "80 8080"
            unset options
            set http-policy enable
            unset post-lang
        end
        ...
    next
end

```

Session-based with web-auth cookie

The web-auth-cookie feature is necessary for session-based authentication under transparent web-proxy.

The configuration is the same as for IP-based authentication, except `ip-based` is disabled in the authentication rule:

```

config authentication rule
    edit "kerberos-rules"
        set status <enable>
        set protocol <http>
        set srcaddr "all"
        set ip-based <disable>
        set active-auth-method <kerberos-scheme>
    next

config authentication setting
    set captive-portal <fqdn-name>
    set captive-portal-port "9998"
end

```

HTTP tunnel authentication

You can trigger user authentication on HTTP CONNECT request at the policy level. A new CLI entry has been added under `config firewall proxy-policy` which will trigger the authentication process `get-user`, even when there is no user or group configured.

Note that, as shown below, explicit web proxy must be set.

Syntax

```
config firewall proxy-policy
  edit {policyid}
    set proxy explicit-web
    set http-tunnel-auth {enable | disable}
  next
end
```

Certificate authentication

You can configure certificate-based authentication for FortiGate administrators, SSL VPN users, and IPsec VPN users. See [Configuring certificate-based authentication on page 255](#).

Certificates are also inherent to the HTTPS protocol, where the browser validates the server's identity using certificates. A site certificate must be installed on the FortiGate unit and the corresponding Certificate Authority (CA) certificate installed in the web browser.

To force the use of HTTPS, go to **User & Device > Authentication Settings** and select **Redirect HTTP Challenge to a Secure Channel (HTTPS)**.

Restricting number of concurrent user logins

Some users on your network may often have multiple account sessions open at one time either to the same network resource or accessing to the admin interface on the FortiGate unit.

While there are valid reasons for having multiple concurrent sessions open, hackers also do this to speed up their malicious work. Often a hacker is making multiple attempts to gain access to the internal network or the admin interface of the FortiGate unit, usually from different IP addresses to appear to the FortiGate unit as legitimate users. For this reason, the more concurrent sessions a hacker has open at once, the faster they will achieve their goal.

To help prevent this, you can disallow concurrent administrative access using the same administrator user name. This allows only one session with the same username even if it is from the same IP.

To disable concurrent administrator sessions - CLI:

```
config system global
  set admin-concurrent disable
end
```

VPN authentication

All VPN configurations require users to authenticate. Authentication based on user groups applies to:

- SSL VPNs
- PPTP and L2TP VPNs
- an IPsec VPN that authenticates users using dialup groups
- a dialup IPsec VPN that uses XAUTH authentication (Phase 1)

You must create user accounts and user groups before performing the procedures in this section. If you create a user group for dialup IPsec clients or peers that have unique peer IDs, their user accounts must be stored locally on the FortiGate unit. You cannot authenticate these types of users using a RADIUS or LDAP server.

Configuring authentication of SSL VPN users

The general procedure for authenticating SSL VPN users is:

1. Configure user accounts.
2. Create one or more user groups for SSL VPN users.
3. Enable SSL VPN.
4. Optionally, set inactivity and authentication timeouts.
5. Configure a security policy with the user groups you created for SSL VPN users.
See FortiOS Handbook SSL VPN guide.

Configuring authentication timeout

By default, the SSL VPN authentication expires after 8 hours (28 800 seconds). You can change it only in the CLI, and the time entered must be in seconds. The maximum time is 72 hours (259 200 seconds). For example, to change this timeout to one hour, you would enter:

```
config vpn ssl settings
    set auth-timeout 3600
end
```

If you set the authentication timeout (`auth-timeout`) to 0 when you configure the timeout settings, the remote client does not have to re-authenticate unless they log out of the system. To fully take advantage of this setting, the value for `idle-timeout` has to be set to 0 also, so that the client does not time out if the maximum idle time is reached. If the `idle-timeout` is not set to the infinite value, the system will log out if it reaches the limit set, regardless of the `auth-timeout` setting.

Configuring authentication of remote IPsec VPN users

An IPsec VPN on a FortiGate unit can authenticate remote users through a dialup group. The user account name is the peer ID and the password is the pre-shared key.

Authentication through user groups is supported for groups containing only local users. To authenticate users using a RADIUS or LDAP server, you must configure XAUTH settings. See Configuring XAuth authentication.

To configure user group authentication for dialup IPsec - web-based manager:

1. Configure the dialup users who are permitted to use this VPN. Create a user group with **Type** set to **Firewall** and add them to it.
For more information, see [Users and user groups on page 181](#)
2. Go to **VPN > IPsec Wizard**, select **Remote Access**, choose a name for the VPN, and enter the following information.

Incoming Interface	Select the incoming interface name.
Authentication Method	List of authentication methods available for users. Select Pre-shared Key and enter the pre-shared key.
User Group	Select the user group that is to be allowed access to the VPN. The listed user groups contain only users with passwords on the FortiGate unit.

3. Select **Next** and continue configure other VPN parameters as needed.
4. Select **OK**.

To configure user group authentication for dialup IPsec - CLI example:

The `peertype` and `usrgrp` options configure user group-based authentication.

```
config vpn ipsec phase1
edit office_vpn
set interface port1
set type dynamic
set psksecret yORRAzltNGhzgtV32jend
set proposal 3des-sha1 aes128-sha1
set peertype dialup
set usrgrp Group1
end
```

Configuring XAuth authentication

Extended Authentication (XAuth) increases security by requiring additional user authentication information in a separate exchange at the end of the VPN Phase 1 negotiation. The FortiGate unit asks the user for a username and password. It then forwards the user's credentials (the password is encrypted) to an external RADIUS or LDAP server for verification.

XAuth can be used in addition to or in place of IPsec phase 1 peer options to provide access security through an LDAP or RADIUS authentication server. You must configure a dialup user group whose members are all externally authenticated.

To configure authentication for a dialup IPsec VPN - web-based manager:

1. Configure the users who are permitted to use this VPN. Create a user group and add the users to the group. For more information, see ["Users and user groups" on page 181](#).
2. Go to **VPN > IPsec Wizard**, select **Remote Access**, choose a name for the VPN, and enter the following information.

Incoming Interface	Select the incoming interface name.
Authentication Method	List of authentication methods available for users. Select Pre-shared Key and enter the pre-shared key.
User Group	Select the user group that is to be allowed access to the VPN. The listed user groups contain only users with passwords on the FortiGate unit.

3. Select **Next** and continue configure other VPN parameters as needed.
4. Select **OK**.
5. Go to **VPN > IPsec Tunnels**, edit the Tunnel just created, select **Convert To Custom Tunnel**, and edit **XAUTH** as following:

Type	Select PAP , CHAP , or AUTO . Use CHAP whenever possible. Use PAP with all implementations of LDAP and with other authentication servers that do not support CHAP, including some implementations of Microsoft RADIUS. Use AUTO with the Fortinet Remote VPN Client and where the authentication server supports CHAP but the XAuth client does not.
-------------	---

User Group

Select the user group that is to have access to the VPN. The list of user groups does not include any group that has members whose password is stored on the FortiGate unit.

6. Select OK.

For more information about XAUTH configuration, see the IPsec VPN chapter of the FortiOS Handbook.

To configure authentication for a dialup IPsec VPN - CLI example:

The `xauthtype` and `authusrgrp` fields configure XAuth authentication.

```
config vpn ipsec phase1
  edit office_vpn
    set interface port1
    set type dynamic
    set psksecret yORRAzltNGhzgtV32jend
    set proposal 3des-sha1 aes128-sha1
    set peertype dialup
    set xauthtype pap
    set usrgrp Group1
  end
```

Some parameters specific to setting up the VPN itself are not shown here. For detailed information about configuring IPsec VPNs, see the FortiOS Handbook IPsec VPN guide.

Configuring authentication of PPTP VPN users and user groups

Configuration of a PPTP VPN is possible only through the CLI. You can configure user groups and security policies using either CLI or web-based manager.



LDAP user authentication is supported for PPTP, L2TP, IPsec VPN, and firewall authentication.

However, with PPTP, L2TP, and IPsec VPN, PAP (Packet Authentication Protocol) is supported, while CHAP (Challenge Handshake Authentication Protocol) is not.

To configure authentication for a PPTP VPN

1. Configure the users who are permitted to use this VPN. Create a security user group and add them to it. For more information, see [Users and user groups on page 181](#).
2. Configure the PPTP VPN in the CLI as in this example.

```
config vpn pptp
  set status enable
  set sip 192.168.0.100
  set eip 192.168.0.110
  set usrgrp PPTP_Group
end
```

The `sip` and `eip` fields define a range of virtual IP addresses assigned to PPTP clients.

Configure a security policy. The source interface is the one through which the clients will connect. The source address is the PPTP virtual IP address range. The destination interface and address depend on the network to which the clients will connect. The policy action is ACCEPT.

Configuring authentication of L2TP VPN users/user groups

Configuration of a L2TP VPN is possible only through the CLI. You can configure user groups and security policies using either CLI or web-based manager.



LDAP user authentication is supported for PPTP, L2TP, IPsec VPN, and firewall authentication.

However, with PPTP, L2TP, and IPsec VPN, PAP (Packet Authentication Protocol) is supported, while CHAP (Challenge Handshake Authentication Protocol) is not.

To configure authentication for a L2TP VPN

1. Configure the users who are permitted to use this VPN. Create a user group and add them to it. For more information, see [Users and user groups on page 181](#).
2. Configure the L2TP VPN in the CLI as in this example.

```
config vpn l2tp
  set status enable
  set sip 192.168.0.100
  set eip 192.168.0.110
  set usrgroup L2TP_Group
end
```

The `sip` and `eip` fields define a range of virtual IP addresses assigned to L2TP clients.

3. Configure a security policy. The source interface is the one through which the clients will connect. The source address is the L2TP virtual IP address range. The destination interface and address depend on the network to which the clients will connect. The policy action is ACCEPT.

Captive portals

A captive portal is a convenient way to authenticate web users on wired or WiFi networks.

This section describes:

- [Introduction to captive portals](#)
- [Configuring a captive portal](#)
- [Customizing captive portal pages](#)

Introduction to captive portals

You can authenticate your users on a web page that requests the user's name and password. Until the user authenticates successfully, the authentication page is returned in response to any HTTP request. This is called a captive portal.

After successful authentication, the user accesses the requested URL and can access other web resources, as permitted by security policies. Optionally, the captive portal itself can allow web access to only the members of specified user group.

The captive portal can be hosted on the FortiGate unit or on an external authentication server. You can configure captive portal authentication on any network interface, including WiFi and VLAN interfaces.

When a captive portal is configured on a WiFi interface, the access point initially appears open. The wireless client can connect to the access point with no security credentials, but sees only the captive portal authentication page.

WiFi captive portal types:

- **Authentication** — until the user enters valid credentials, no communication beyond the AP is permitted.
- **Disclaimer + Authentication** — immediately after successful authentication, the portal presents the disclaimer page—an acceptable use policy or other legal statement—to which the user must agree before proceeding.
- **Disclaimer Only** — the portal presents the disclaimer page—an acceptable use policy or other legal statement—to which the user must agree before proceeding. The authentication page is not presented.
- **Email Collection** — the portal presents a page requesting the user's email address, for the purpose of contacting the person in future. This is often used by businesses who provide free WiFi access to their customers. The authentication page is not presented.

Configuring a captive portal

Captive portals are configured on network interfaces. On a physical (wired) network interface, you edit the interface configuration in **Network > Interfaces** and set **Security Mode** to **Captive Portal**. A WiFi interface does not exist until the WiFi SSID is created. You can configure a WiFi captive portal at the time that you create the SSID. Afterwards, the captive portal settings will also be available by editing the WiFi network interface in **Network > Interfaces**.

To configure a wired Captive Portal - web-based manager:

1. Go to **Network > Interfaces** and edit the interface to which the users connect.
2. In **Security Mode** select **Captive Portal**.

Security Mode	Captive Portal ▼
Authentication Portal	<input checked="" type="radio"/> Local <input type="radio"/> External
User Groups	Use Groups from Policies ▼
Exempt List	Click to set... ▼
Customize Portal Messages	<input type="checkbox"/>

3. Enter

Authentication Portal	Local - portal hosted on the FortiGate unit. Remote - enter FQDN or IP address of external portal.
User Groups	Select permitted user groups or select Use Groups from Policies , which permits the groups specified in the security policy. Use Groups from Policies is not available in WiFi captive portals.
Exempt List	Select exempt lists whose members will not be subject to captive portal authentication.
Customize Portal Messages	Enable, then select Edit. See Customizing captive portal pages on page 235 .

4. Select **OK**.**To configure a WiFi captive portal - web-based manager:**

- Go to **WiFi & Switch Controller > SSID** and create your SSID.
If the SSID already exists, you can edit the SSID or you can edit the WiFi interface in **Network > Interfaces**.
- In **Security Mode**, select **Captive Portal**.

WiFi Settings	
SSID	fortinet
Security Mode	Captive Portal ▼
Client Limit	<input type="checkbox"/>
Portal Type	<input checked="" type="radio"/> Authentication <input type="radio"/> Disclaimer + Authentication <input type="radio"/> Disclaimer Only <input type="radio"/> Email Collection
Authentication Portal	<input type="radio"/> Local <input checked="" type="radio"/> External example.com/captive/
User Groups	<input type="text" value="+"/> +
Exempt Sources	<input type="text" value="+"/> +
Exempt Destinations/Services	<input type="text" value="+"/> +
Redirect after Captive Portal	<input checked="" type="radio"/> Original Request <input type="radio"/> Specific URL

3. Enter

Portal Type	The portal can provide authentication and/or disclaimer, or perform user email address collection. See Introduction to captive portals on page 233 .
Authentication Portal	Local - portal hosted on the FortiGate unit. Remote - enter FQDN or IP address of external portal.
User Groups	Select permitted user groups.
Exempt List	Select exempt lists whose members will not be subject to captive portal authentication.
Customize Portal Messages	Click the link of the portal page that you want to modify. See "Captive portals" on page 235 .

4. Select **OK**.

Exemption from the captive portal

A captive portal requires all users on the interface to authenticate. But some devices are not able to authenticate. You can create an exemption list of these devices. For example, a printer might need to access the Internet for firmware upgrades. Using the CLI, you can create an exemption list to exempt all printers from authentication.

```
config user security-exempt-list
  edit r_exempt
    config rule
      edit 1
        set devices printer
      end
    end
  end
```

IPv6 and captive portals

Captive portal supports IPv6. Host name and address commands are available under `config auth setting`:

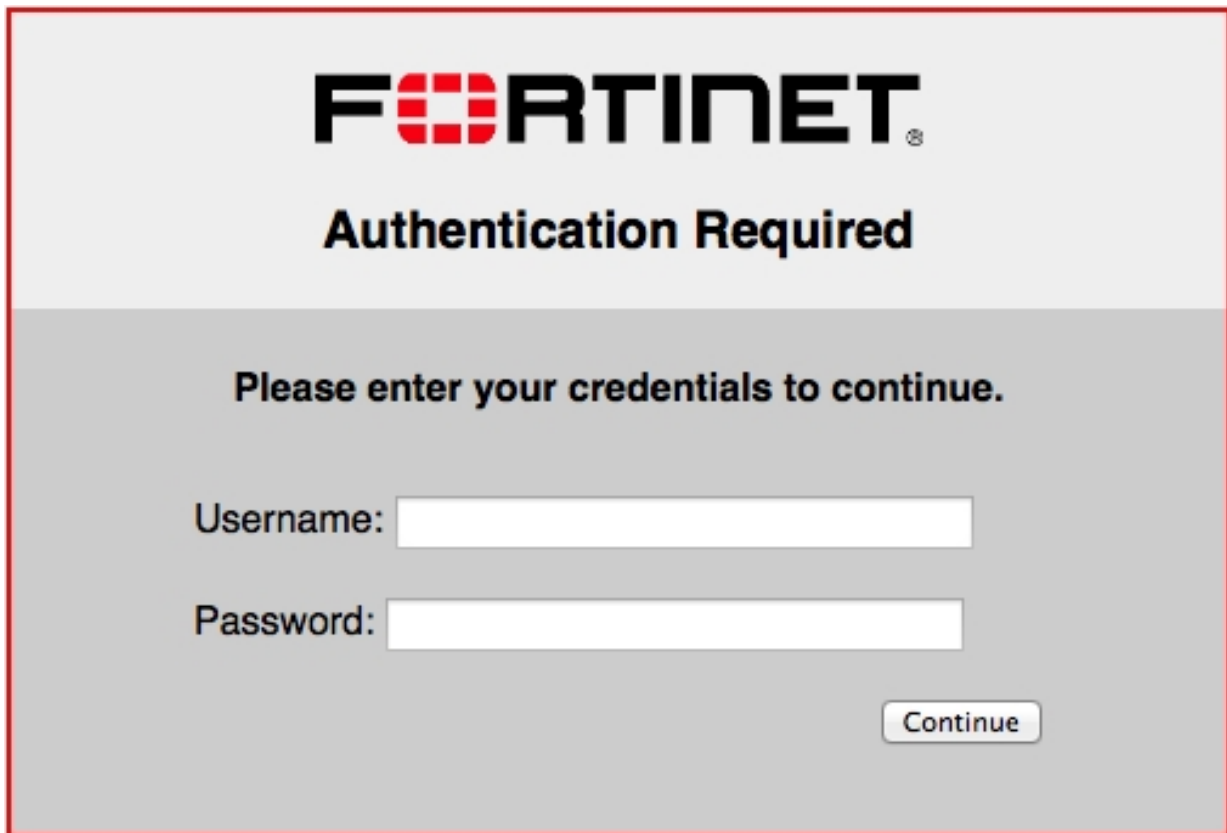
```
config auth setting
  set captive-portal6      --> IPv6 captive portal host name
  set captive-portal-ip6   --> Captive portal IPv6 address
end
```

Customizing captive portal pages

These pages are defined in replacement messages. Defaults are provided. In the web-based manager, you can modify the default messages in the SSID configuration by selecting **Customize Portal Messages**. Each SSID can have its own unique portal content.

The captive portal contains the following default web pages:

- **Login page**—requests user credentials

The image shows a Fortinet captive portal login page. At the top, the Fortinet logo is displayed in black and red. Below the logo, the text "Authentication Required" is centered in a bold, black font. Underneath this, a message "Please enter your credentials to continue." is centered. There are two input fields: "Username:" followed by a white text box, and "Password:" followed by a white text box. A "Continue" button is located at the bottom right of the form area.

Typical modifications for this page would be to change the logo and modify some of the text.

You can change any text that is not part of the HTML code nor a special tag enclosed in double percent (%) characters.

There is an exception to this rule. The line "Please enter your credentials to continue" is provided by the `%%QUESTION%%` tag. You can replace this tag with text of your choice. Except for this item, you should not remove any tags because they may carry information that the FortiGate unit needs.

- **Login failed page**—reports that the entered credentials were incorrect and enables the user to try again.

The image shows a web page for Fortinet authentication. At the top, the Fortinet logo is displayed in black and red. Below the logo, the text "Authentication Failed" is centered in a bold, black font. Underneath this, a message reads "Firewall Authentication failed. Please try again." in a bold, black font. Below the message, there are two input fields: "Username:" followed by a white text box, and "Password:" followed by a white text box. At the bottom right of the form, there is a button labeled "Continue". The entire form is enclosed in a red border.

The Login failed page is similar to the Login page. It even contains the same login form. You can change any text that is not part of the HTML code nor a special tag enclosed in double percent (%) characters.

There is an exception to this rule. The line “Firewall authentication failed. Please try again.” is provided by the `%%FAILED_MESSAGE%%` tag. You can replace this tag with text of your choice. Except for this item, you should not remove any tags because they may carry information that the FortiGate unit needs.

- **Disclaimer page**—is a statement of the legal responsibilities of the user and the host organization to which the user must agree before proceeding. (WiFi or SSL VPN only)

Terms and Disclaimer Agreement

You are about to access Internet content that is not under the control of the network access provider. The network access provider is therefore not responsible for any of these sites, their content or their privacy policies. The network access provider and its staff do not endorse nor make any representations about these sites, or any information, software or other products or materials found there, or any results that may be obtained from using them. If you decide to access any Internet content, you do this entirely at your own risk and you are responsible for ensuring that any accessed material does not infringe the laws governing, but not exhaustively covering, copyright, trademarks, pornography, or any other material which is slanderous, defamatory or might cause offence in any other way.

Do you agree to the above terms?

- **Declined disclaimer page**—is displayed if the user does not agree to the statement on the Disclaimer page. Access is denied until the user agrees to the disclaimer.



Changing images in portal messages

You can replace the default Fortinet logo with your organization's logo. First, import the logo file into the FortiGate unit and then modify the Login page code to reference your file.

To import a logo file:

1. Go to **System > Replacement Messages** and select **Manage Images**.
2. Select **Create New**.
3. Enter a **Name** for the logo and select the appropriate **Content Type**.
The file must not exceed 24 Kilo bytes.
4. Select **Browse**, find your logo file and then select **Open**.
5. Select **OK**.

To specify the new logo in the replacement message:

1. Go to **Network > Interfaces** and edit the interface.
The **Security Mode** must be **Captive Portal**.
2. Select the portal message to edit.
 - In SSL VPN or WiFi interfaces, in **Customize Portal Messages** click the link to the portal messages that you want to edit.
 - In other interfaces, make sure that **Customize Portal Messages** is selected, select the adjacent **Edit** icon, then select the message that you want to edit.
3. In the HTML message text, find the %%IMAGE tag.
By default it specifies the Fortinet logo: %%IMAGE:logo_fw_auth%%
4. Change the image name to the one you provided for your logo.
The tag should now read, for example, %%IMAGE:mylogo%%
5. Select **Save**.
6. Select **OK**.

Modifying text in portal messages

Generally, you can change any text that is not part of the HTML code nor a special tag enclosed in double percent (%) characters. You should not remove any tags because they may carry information that the FortiGate unit needs. See the preceding section for any exceptions to this rule for particular pages.

To modify portal page text

1. Go to **Network > Interfaces** and edit the interface.
The SSID **Security Mode** must be **Captive Portal**.
2. Select the portal message to edit.
 - In SSL VPN or WiFi interfaces, in **Customize Portal Messages** click the link to the portal messages that you want to edit.
 - In other interfaces, make sure that **Customize Portal Messages** is selected, select the adjacent **Edit** icon, then select the message that you want to edit.
3. Edit the HTML message text, then select **Save**.
4. Select **OK**.

Configuring disclaimer page for ethernet interface captive portals

While you can customize a disclaimer page for captive portals that connect via WiFi, the same can be done for wired connections. However, this can only be configured on the CLI Console, and only without configuring user groups.

When configuring a captive portal through the CLI, you may set `security-groups` to a specific user group. The result of this configuration will show an authentication form to users who wish to log in to the captive portal—**not** a disclaimer page. If you do not set any `security-groups` in your configuration, an "Allow all" status will be in effect, and the disclaimer page will be displayed for users.

The example CLI configuration below shows setting up a captive portal interface without setting security-groups, resulting in a disclaimer page for users:

```
config system interface
  edit "port1"
    set vdom "root"
    set ip 172.16.101.1 255.255.255.0
    set allowaccess ping https ssh snmp http
    set type physical
    set explicit-web-proxy enable
    set alias "LAN"
    set security-mode captive-portal
    set snmp-index 1
  next
end
```

Certificate-based authentication

This section provides an overview of how the FortiGate unit verifies the identities of administrators, SSL VPN users, or IPsec VPN peers using X.509 security certificates.

The following topics are included in this section:

- [What is a security certificate?](#)
- [Certificates overview](#)
- [Managing X.509 certificates](#)
- [Configuring certificate-based authentication](#)
- [Support for per-VDOM certificates](#)
- [Example — Generate a CSR on the FortiGate unit](#)
- [Example — Generate and Import CA certificate with private key pair on OpenSSL](#)
- [Example — Generate an SSL certificate in OpenSSL](#)

What is a security certificate?

A security certificate is a small text file that is part of a third-party generated public key infrastructure (PKI) to help guarantee the identity of both the user logging on and the web site they where they are logging in.

A certificate includes identifying information such as the company and location information for the web site, as well as the third-party company name, the expiry date of the certificate, and the public key.

FortiGate units use X.509 certificates to authenticate single sign-on (SSO) for users. The X.509 standard has been in use since before 2000, but has gained popularity with the Internet's increased popularity. X.509 v3 is defined in RFC 5280 and specifies standard formats for public key certificates, certificate revocation lists, and a certification path validation algorithm. The unused earlier X.509 version 1 was defined in RFC 1422.

The main difference between X.509 and PGP certificates is that where in PGP anyone can sign a certificate, for X.509 only a trusted authority can sign certificates. This limits the source of certificates to well known and trustworthy sources. Where PGP is well suited for one-to-one communications, the X.509 infrastructure is intended to be used in many different situations including one-to-many communications. Some common filename extensions for X.509 certificates are listed below.

Common certificate filename extensions

Filetype	Format name	Description
.pem	Privacy Enhanced Mail (PEM)	Base64 encoded DER certificate, that uses: “-----BEGIN CERTIFICATE-----” and “-----END CERTIFICATE-----”
.cer .crt .der	Security Certificate	Usually binary DER form, but Base64-encoded certificates are common too.
.p7b .p7c		Structure without data, just certificates or CRLs. PKCS#7 is a standard for signing or encrypting (officially called “enveloping”) data.
.p12	PKCS#12	May contain certificate(s) (public) and private keys (password protected).
.pfx	personal information exchange (PFX)	Older format. Came before PKCS#12. Usually today data is in PKCS#12 format.

Certificates overview

Certificates play a major role in authentication of clients connecting to network services via HTTPS, both for administrators and SSL VPN users. Certificate authentication is optional for IPsec VPN peers.

This section includes:

- [Certificates and protocols](#)
- [IPsec VPNs and certificates](#)
- [Certificate types on the FortiGate unit](#)



Public CA certificates found on the FortiGate are provided through firmware upgrades and installations.

Certificates and protocols

There are a number of protocols that are commonly used with certificates including SSL and HTTPS, and other certificate-related protocols.

SSL and HTTPS

The secure HTTP (HTTPS) protocol uses SSL. Certificates are an integral part of SSL. When a web browser connects to the FortiGate unit via HTTPS, a certificate is used to verify the FortiGate unit's identity to the client. Optionally, the FortiGate unit can require the client to authenticate itself in return.

By default, the FortiGate unit uses a self-signed security certificate to authenticate itself to HTTPS clients. When the certificate is offered, the client browser displays two security messages.

- The first message prompts users to accept and optionally install the FortiGate unit's self-signed security certificate. If the user does not accept the certificate, the FortiGate unit refuses the connection. When the user accepts the certificate, the FortiGate login page is displayed, and the credentials entered by the user are encrypted before they are sent to the FortiGate unit. If the user chooses to install the certificate, the prompt is not displayed again.
- Just before the FortiGate login page is displayed, a second message informs users that the FortiGate certificate distinguished name differs from the original request. This message is displayed because the FortiGate unit redirects the connection (away from the distinguished name recorded in the self-signed certificate) and can be ignored.

Optionally, you can install an X.509 server certificate issued by a certificate authority (CA) on the FortiGate unit. You can then configure the FortiGate unit to identify itself using the server certificate instead of the self-signed certificate.

For more information, see the FortiOS Handbook SSL VPN guide.

After successful certificate authentication, communication between the client browser and the FortiGate unit is encrypted using SSL over the HTTPS link.

Certificate-related protocols

There are multiple protocols that are required for handling certificates. These include the Online Certificate Status Protocol (OCSP), Simple Certificate Enrollment Protocol (SCEP), Server-based Certificate Validation Protocol (SCVP), and Certificate Management Protocol (CMP).

Online Certificate Status Protocol

Online Certificate Status Protocol (OCSP) allows the verification of X.509 certificate expiration dates. This is important to prevent hackers from changing the expiry date on an old certificate to a future date.

Normally certificate revocation lists (CRLs) are used, but OCSP is an alternate method available. However a CRL is a public list, and some companies may want to avoid the public exposure of their certificate structure even if it is only invalid certificates.

The OSCP check on the certificate's revocation status is typically carried out over HTTP with a request-response format. The authority responding can reply with a status of good, revoked, or unknown for the certificate in question.

Simple Certificate Enrollment Protocol

Simple Certificate Enrollment Protocol (SCEP) is an automated method of signing up for certificates. Typically this involves generating a request you send directly to the SCEP service, instead of generating a file request that may or may not be signed locally.

Server-based Certificate Validation Protocol

Server-based Certificate Validation Protocol (SCVP) is used to trace a certificate back to a valid root level certificate. This ensures that each step along the path is valid and trustworthy.

Certificate Management Protocol version 2

Certificate Management Protocol version 2 (CMPv2) is an enrollment and revocation protocol for certificates.

IPsec VPNs and certificates

Certificate authentication is a more secure alternative to pre-shared key (shared secret) authentication for IPsec VPN peers. Unlike administrators or SSL VPN users, IPsec peers use HTTP to connect to the VPN gateway configured on the FortiGate unit. The VPN gateway configuration can require certificate authentication before it permits an IPsec tunnel to be established. See [Authenticating IPsec VPN users with security certificates on page 256](#).

Certificate types on the FortiGate unit

There are different types of certificates available that vary depending on their intended use. FortiOS supports local, remote, CA, and CRL certificates.

Local certificates

Local certificates are issued for a specific server, or web site. Generally they are very specific, and often for an internal enterprise network. For example a personal web site for John Smith at www.example.com (such as <http://www.example.com/home/jsmith>) would have its own local certificate.

These can optionally be just the certificate file, or also include a private key file and PEM passphrase for added security.

For information about generating a certificate request, see [Generating a certificate signing request on page 246](#). For information about installing a local certificate, see [Obtaining and installing a signed server certificate from an external CA on page 249](#).

Remote certificates

Remote certificates are public certificates without a private key. For dynamic certificate revocation, you need to use an Online Certificate Status Protocol (OCSP) server. The OCSP is configured in the CLI only. Installed Remote (OCSP) certificates are displayed in the Remote Certificates list. You can select **Import** to install a certificate from the management PC.

CA root certificates

CA root certificates are similar to local certificates, however they apply to a broader range of addresses or to whole company; they are one step higher up in the organizational chain. Using the local certificate example, a CA root certificate would be issued for all of www.example.com instead of just the smaller single web page.

Certificate revocation list

Certificate revocation list (CRL) is a list of certificates that have been revoked and are no longer usable. This list includes certificates that have expired, been stolen, or otherwise compromised. If your certificate is on this list, it will not be accepted. CRLs are maintained by the CA that issues the certificates and includes the date and time when the next CRL will be issued as well as a sequence number to help ensure you have the most current version of the CRL.

Certificate signing

The trust in a certificate comes from the authority that signs it. For example if VeriSign signs your CA root certificate, it is trusted by everyone. While these certificates are universally accepted, it is cumbersome and expensive to have all certificates on a corporate network signed with this level of trust.

With self-signed certificates nobody, except the other end of your communication, knows who you are and therefore they do not trust you as an authority. However this level is useful for encryption between two points — neither point may care about who signed the certificate, just that it allows both points to communicate. This is very useful for internal networks and communications.

A general rule is that CA signed certificates are accepted and sometimes required, but it is easier to self-sign certificates when you are able.

For more on the methods of certificate signing see [Generating a certificate signing request on page 246](#).

BIOS certificate compatibility

FortiOS supports backwards compatibility between BIOS version 4 and BIOS version 3.

BIOS V4 certificates:

- Fortinet_CA
- Fortinet_Sub_CA
- Fortinet_Factory

BIOS V3 certificates:

- Fortinet_CA_Backup
- Fortinet_Factory_Backup

When FortiOS connects to FortiGuard, FortiCloud, FortiManager, FortiAnalyzer, FortiSandbox as a client, the BIOS certificate **Fortinet_Factory** will be the default client certificate. When the server returns its certificate (chain) back, FortiOS looks up the issuer of the server certificate and either keeps client certificate as is or switches to the BIOS certificate **Fortinet_Factory_Backup**. This process occurs in one handshake.

When FortiOS connects to FortiCare, the BIOS certificate **Fortinet_Factory** is the only client certificate and Server Name Indication (SNI) is set. There is no switchover of certificate during SSL handshake.

When FortiOS acts as a server when connected by FortiExtender, FortiSwitch, FortiAP, etc., **Fortinet_Factory** is the default server certificate. FortiOS detects SNI in client hello, and if no SNI is found or if the CN in SNI is different from the CN of **Fortinet_CA**, it switches to use the **Fortinet_Factory_Backup**.

Managing X.509 certificates

Managing security certificates is required due to the number of steps involved in both having a certificate request signed, and then distributing the correct files for use.

You use the FortiGate unit or CA software such as OpenSSL to generate a certificate request. That request is a text file that you send to the CA for verification, or alternately you use CA software to self-validate. Once validated, the certificate file is generated and must be imported to the FortiGate unit before it can be used. These steps are explained in more detail later in this section.

This section provides procedures for generating certificate requests, installing signed server certificates, and importing CA root certificates and CRLs to the FortiGate unit.

For information about how to install root certificates, CRLs, and personal or group certificates on a remote client browser, refer to your browser's documentation.



As a requirement of Network Device Collaborative Protection Profile (NDcPP), FortiOS supports and handles the use of wildcards for the following certificate reference parameters:

- Subject Alternative Name (SAN)
- Common Name (CN)

This section includes:

- [Generating a certificate signing request](#)
- [Generating certificates with CA software](#)
- [Obtaining and installing a signed server certificate from an external CA](#)
- [Installing a CA root certificate and CRL to authenticate remote clients](#)
- [ExtendedKeyUsage for x.509 certificates](#)

Generating a certificate signing request

Whether you create certificates locally with a software application or obtain them from an external certificate service, you will need to generate a certificate signing request (CSR).

When you generate a CSR, a private and public key pair is created for the FortiGate unit. The generated request includes the public key of the FortiGate unit and information such as the FortiGate unit's public static IP address, domain name, or email address. The FortiGate unit's private key remains confidential on the FortiGate unit.

After you submit the request to a CA, the CA will verify the information and register the contact information on a digital certificate that contains a serial number, an expiration date, and the public key of the CA. The CA will then sign the certificate, and you install the certificate on the FortiGate unit.

The Certificate Request Standard is a public key cryptography standard (PKCS) published by RSA, specifically PKCS10 which defines the format for CSRs. This is defined in RFC 2986.

To generate a certificate request in FortiOS - web-based manager:

1. Go to **System > Certificates**.
2. Select **Generate**.
3. In the **Certificate Name** field, enter a unique meaningful name for the certificate request. Typically, this would be the hostname or serial number of the FortiGate unit or the domain of the FortiGate unit such as example.com.



Do not include spaces in the certificate name. This will ensure compatibility of a signed certificate as a PKCS12 file to be exported later on if required.



Prior to FortiOS 5.4, passwords for local certificates that were generated via either SCEP or CLI could not have their passwords reset. Passwords can be set in the CLI using the following command:

```
config vpn certificate local
edit <name>
set password <password>
next
end
```

4. Enter values in the **Subject Information** area to identify the FortiGate unit:
 - If the FortiGate unit has a static IP address, select **Host IP** and enter the public IP address of the FortiGate unit. If the FortiGate unit does not have a public IP address, use an email address (or fully qualified domain name (FQDN) if available) instead.
 - If the FortiGate unit has a dynamic IP address and subscribes to a dynamic DNS service, use a FQDN if available to identify the FortiGate unit. If you select **Domain Name**, enter the FQDN of the FortiGate unit. Do not include the protocol specification (http://) or any port number or path names.



If a domain name is not available and the FortiGate unit subscribes to a dynamic DNS service, an “unable to verify certificate” type message may be displayed in the user’s browser whenever the public IP address of the FortiGate unit changes.

- If you select **E-Mail**, enter the email address of the owner of the FortiGate unit.
5. Enter values in the **Optional Information** area to further identify the FortiGate unit.

Organization Unit	Name of your department. You can enter a series of OUs up to a maximum of 5. To add or remove an OU, use the plus (+) or minus (-) icon.
Organization	Legal name of your company or organization.
Locality (City)	Name of the city or town where the FortiGate unit is installed.
State/Province	Name of the state or province where the FortiGate unit is installed.
Country	Select the country where the FortiGate unit is installed.
e-mail	Contact email address.
Subject Alternative Name	<p>Optionally, enter one or more alternative names for which the certificate is also valid. Separate names with a comma. A name can be:</p> <ul style="list-style-type: none"> • e-mail address • IP address • URI • DNS name (alternatives to the Common Name) • directory name (alternatives to the Distinguished Name) <p>You must precede the name with the name type. Examples:</p> <p>IP: 1.1.1.1 email: test@fortinet.com email: my@other.address URI: http://my.url.here/</p>
Password for private key	Option to export local certificate and its private key in password protected p12.

6. From the **Key Type** list, select **RSA** or **Elliptic Curve**.
7. From the **Key Size** list, select **1024 Bit**, **1536 Bit**, **2048 Bit**, **4096 Bit** or **secp256r1**, **secp384r1**, **secp521r1** respectively. Larger keys are slower to generate but more secure.

8. In **Enrollment Method**, you have two methods to choose from. Select **File Based** to generate the certificate request, or **Online SCEP** to obtain a signed SCEP-based certificate automatically over the network. For the SCEP method, enter the URL of the SCEP server from which to retrieve the CA certificate, and the CA server challenge password.
 9. Select **OK**.
 10. The request is generated and displayed in the **Local Certificates** list with a status of `PENDING`.
 11. Select the **Download** button to download the request to the management computer.
 12. In the **File Download** dialog box, select **Save** and save the Certificate Signing Request on the local file system of the management computer.
 13. Name the file and save it on the local file system of the management computer.
- The certificate request is ready for the certificate authority to be signed.

Generating certificates with CA software

CA software allows you to generate unmanaged certificates and CA certificates for managing other certificates locally without using an external CA service. Examples of CA software include `ssl-ca` from OpenSSL (available for Linux, Windows, and Mac) or `gensslcert` from SuSE, MS Windows Server 2000 and 2003 come with a CA as part of their certificate services, and in MS Windows 2008 CA software can be installed as part of the Active Directory installation. See [Example — Generate and Import CA certificate with private key pair on OpenSSL on page 259](#).

The general steps for generating certificates with CA software are

1. Install the CA software as a stand-alone root CA.
2. Provide identifying information for your self-administered CA.

While following these steps, the methods vary slightly when generating server certificates, CA certificates, and PKI certificates.

Server certificate

1. Generate a Certificate Signing Request (CSR) on the FortiGate unit.
2. Copy the CSR base-64 encoded text (PKCS10 or PKCS7) into the CA software and generate the certificate. PKCS10 is the format used to send the certificate request to the signing authority. PKCS7 is the format the signing authority can use for the newly signed certificate.
3. Export the certificate as a X.509 DER encoded binary file with `.CER` extension
4. Upload the certificate file to the FortiGate unit Local Certificates page (type is Certificate).

CA certificate

1. Retrieve the CA Certificate from the CA software as a DER encoded file.
2. Import the CA certificate file to the FortiGate unit at **System > Certificates** and select **Import > Certificates**.

PKI certificate

1. Generate a Certificate Signing Request (CSR) on the FortiGate unit.
2. Copy the CSR base-64 encoded text (PKCS#10 or PKCS#7) into the CA software and generate the certificate. PKCS10 is the format used to send the certificate request to the signing authority. PKCS7 is the format the signing authority can use for the newly signed certificate.
3. Export the certificate as a X.509 DER encoded binary file with `.CER` extension.
4. Install the certificate in the user's web browser or IPsec VPN client as needed.

Obtaining and installing a signed server certificate from an external CA

To obtain a signed server certificate for a FortiGate unit, you must send a request to a CA that provides digital certificates that adhere to the X.509 standard. The FortiGate unit provides a way for you to generate the request.

To submit the certificate signing request (file-based enrollment):

1. Using the web browser on the management computer, browse to the CA web site.
2. Follow the CA instructions for a base-64 encoded PKCS#10 certificate request and upload your certificate request.
3. Follow the CA instructions to download their root certificate and CRL.
When you receive the signed server certificate from the CA, install the certificate on the FortiGate unit.

To install or import the signed server certificate - web-based manager

1. On the FortiGate unit, go to **System > Certificates** and select **Import > Local Certificates**.
2. From **Type**, select **Local Certificate**.
3. Select **Browse**, browse to the location on the management computer where the certificate was saved, select the certificate, and then select **Open**.
4. Select **OK**, and then select **Return**.

Installing a CA root certificate and CRL to authenticate remote clients

When you apply for a signed personal or group certificate to install on remote clients, you can obtain the corresponding root certificate and CRL from the issuing CA. When you receive the signed personal or group certificate, install the signed certificate on the remote client(s) according to the browser documentation. Install the corresponding root certificate (and CRL) from the issuing CA on the FortiGate unit according to the procedures given below.

To install a CA root certificate

1. After you download the root certificate of the CA, save the certificate on the management computer. Or, you can use online SCEP to retrieve the certificate.
2. On the FortiGate unit, go to **System > Certificates** and select **Import > CA Certificates**.
3. Do one of the following:
 - To import using SCEP, select **SCEP**. Enter the URL of the SCEP server from which to retrieve the CA certificate. Optionally, enter identifying information of the CA, such as the filename.
 - To import from a file, select **Local PC**, then select **Browse** and find the location on the management computer where the certificate has been saved. Select the certificate, and then select **Open**.
5. Select **OK**, and then select **Return**.

The system assigns a unique name to each CA certificate. The names are numbered consecutively (CA_Cert_1, CA_Cert_2, CA_Cert_3, and so on).

To import a certificate revocation list

A Certificate Revocation List (CRL) is a list of the CA certificate subscribers paired with certificate status information. The list contains the revoked certificates and the reason(s) for revocation. It also records the certificate issue dates and the CAs that issued them.

When configured to support SSL VPNs, the FortiGate unit uses the CRL to ensure that the certificates belonging to the CA and remote peers or clients are valid. The CRL has an “effective date” and a “next update” date. The interval is typically 7 days (for Microsoft CA). FortiOS will update the CRL automatically. Also, there is a CLI command to specify an “update-interval” in seconds. Recommendation should be 24 hours (86400 seconds) but depends on company security policy.

1. After you download the CRL from the CA web site, save the CRL on the management computer.
2. Go to **System > Certificates** and select **Import > CRL**.
3. Do one of the following:
 - To import using an HTTP server, select **HTTP** and enter the URL of the HTTP server.
 - To import using an LDAP server see this KB [article](#).
 - To import using an SCEP server, select **SCEP** and select the Local Certificate from the list. Enter the URL of the SCEP server from which the CRL can be retrieved.
 - To import from a file, select **Local PC**, then select **Browse** and find the location on the management computer where the CRL has been saved. Select the CRL and then select **Open**.
5. Select **OK**, and then select **Return**.

To import a PKCS12 certificate from the CLI

The following CLI syntax can be entered to import a local certificate file:

```
execute vpn certificate local import tftp <file name> <tftp ip address> <file type> <Enter  
for 'cer'>|<password for 'p12'>
```

For example:

```
execute vpn certificate local import tftp FGTF-extern.p12 10.1.100.253 p12 123456
```

In addition, the following CLI syntax can be entered to update certificate bundles from an FTP or TFTP server:

```
execute vpn certificate ca import bundle <file-name.pkg> <ftp/tftp-server-ip>
```

ExtendedKeyUsage for x.509 certificates

As per Network Device Collaborative Protection Profile (NDcPP) v1.0 requirements, server certificates used for TLS connections between FortiGate and FortiAnalyzer have the "Server Authentication" and "Client Authentication" extendedKeyUsage fields in FIPS/CC mode.

The following CLI command is available under `log fortianalyzer` setting to allow you to specify the certificate used to communicate with FortiAnalyzer.

CLI syntax

```
config log fortianalyzer setting  
  set certificate <name>  
end
```

Troubleshooting certificates

There are times when there are problems with certificates — a certificate is seen as expired when its not, or it can't be found. Often the problem is with a third party web site, and not FortiOS. However, some problems can be traced back to FortiOS such as DNS or routing issues.

Enable and disable SHA1 algorithm in SSH key exchanges

In order to investigate your security and conduct compliance testing, a global option allows you to enable/disable SHA1 algorithm in SSH key exchange. Note that, the algorithm is enabled by default.

Syntax

```
config system global
  set ssh-key-sha1 {enable | disable}
end
```

Certificate incorrectly reported as expired

Certificates often are issued for a set period of time such as a day or a month, depending on their intended use. This ensures everyone is using up-to-date certificates. It is also more difficult for hackers to steal and use old certificates.

Reasons a certificate may be reported as expired include:

- It really has expired based on the “best before” date in the certificate
- The FortiGate unit clock is not properly set. If the FortiGate clock is fast, it will see a certificate as expired before the expiry date is really here.
- The requesting server clock is not properly set. A valid example is if your certificate is 2 hours from expiring, a server more than two time zones away would see the certificate as expired. Otherwise, if the server’s clock is set wrongly it will also have the same effect.
- The certificate was revoked by the issuer before the expiry date. This may happen if the issuer believes a certificate was either stolen or misused. Its possible it is due to reasons on the issuer’s side, such as a system change or such. In either case it is best to contact the certificate issuer to determine what is happening and why.

A secure connection cannot be completed (certificate cannot be found)

Everyone who uses a browser has encountered a message such as *This connection is untrusted*. Normally when you try to connect securely to a web site, that web site will present its valid certificate to prove their identity is valid. When the web site’s certificate cannot be verified as valid, the message appears stating *This connection is untrusted* or something similar. If you usually connect to this web site without problems, this error could mean that someone is trying to impersonate or hijack the web site, and best practices dictates you not continue.

Reasons a web site’s certificate cannot be validated include:

- The web site uses an unrecognized self-signed certificate. These are not secure because anyone can sign them. If you accept self-signed certificates you do so at your own risk. Best practices dictate that you must confirm the ID of the web site using some other method before you accept the certificate.
- The certificate is valid for a different domain. A certificate is valid for a specific location, domain, or sub-section of a domain such as one certificate for `support.example.com` that is not valid for `marketing.example.com`. If you encounter this problem, contact the webmaster for the web site to inform them of the problem.
- There is a DNS or routing problem. If the web site’s certificate cannot be verified, it will not be accepted. Generally to be verified, your system checks with the third party certificate signing authority to verify the certificate is valid. If you cannot reach that third party due to some DNS or routing error, the certificate will not be verified.
- Firewall is blocking required ports. Ensure that any firewalls between the requesting computer and the web site allow the secure traffic through the firewall. Otherwise a hole must be opened to allow it through. This includes ports such as 443 (HTTPS) and 22 (SSH).

Online updates to certificates and CRLs

If you obtained your local or CA certificate using SCEP, you can configure online renewal of the certificate before it expires. Similarly, you can receive online updates to CRLs.

Local certificates

In the `config vpn certificate local` command, you can specify automatic certificate renewal. The relevant fields are:

<code>scep-url <URL_str></code>	The URL of the SCEP server. This can be HTTP or HTTPS. The following options appear after you add the <code><URL_str></code> .
<code>scep-password <password_str></code>	The password for the SCEP server.
<code>auto-regenerate-days <days_int></code>	How many days before expiry the FortiGate unit requests an updated local certificate. The default is 0, no auto-update.
<code>auto-regenerate-days-warning <days_int></code>	How many days before local certificate expiry the FortiGate generates a warning message. The default is 0, no warning.

In this example, an updated certificate is requested three days before it expires.

```
config vpn certificate local
edit mycert
set scep-url http://scep.example.com/scep
set scep-server-password my_pass_123
set auto-regenerate-days 3
set auto-regenerate-days-warning 2
end
```

CA certificates

In the `config vpn certificate ca` command, you can specify automatic certificate renewal. The relevant fields are:

Variable	Description
<code>scep-url <URL_str></code>	The URL of the SCEP server. This can be HTTP or HTTPS.
<code>auto-update-days <days_int></code>	How many days before expiry the FortiGate unit requests an updated CA certificate. The default is 0, no auto-update.
<code>auto-update-days-warning <days_int></code>	How many days before CA certificate expiry the FortiGate generates a warning message. The default is 0, no warning.

In this example, an updated certificate is requested three days before it expires.

```
config vpn certificate ca
edit mycert
set scep-url http://scep.example.com/scep
set auto-update-days 3
set auto-update-days-warning 2
```

```
end
```

Certificate revocation lists

If you obtained your CRL using SCEP, you can configure online updates to the CRL using the `config vpn certificate crl` command. The relevant fields are:

Variable	Description
<code>http-url <http_url></code>	URL of the server used for automatic CRL certificate updates. This can be HTTP or HTTPS.
<code>scep-cert <scep_certificate></code>	Local certificate used for SCEP communication for CRL auto-update.
<code>scep-url <scep_url></code>	URL of the SCEP CA server used for automatic CRL certificate updates. This can be HTTP or HTTPS.
<code>update-interval <seconds></code>	How frequently, in seconds, the FortiGate unit checks for an updated CRL. Enter 0 to update the CRL only when it expires. Not available for http URLs.
<code>update-vdom <update_vdom></code>	VDOM used to communicate with remote SCEP server for CRL auto-update.

In this example, an updated CRL is requested only when it expires.

```
config vpn certificate crl
  edit cert_crl
    set http-url http://scep.example.com/scep
    set scep-cert my-scep-cert
    set scep-url http://scep.ca.example.com/scep
    set update-interval 0
    set update-vdom root
  end
```

Backing up and restoring local certificates

The FortiGate unit provides a way to export and import a server certificate and the FortiGate unit's personal key through the CLI. If required (to restore the FortiGate unit configuration), you can import the exported file through the **System > Certificates** page of the web-based manager.



As an alternative, you can back up and restore the entire FortiGate configuration through the **System Information** widget on the Dashboard of the web-based manager. Look for **[Backup]** and **[Restore]** in the **System Configuration** row. The backup file is created in a FortiGate-proprietary format.

To export a server certificate and private key - CLI:

This procedure exports a server (local) certificate and private key together as a password protected PKCS12 file. The export file is created through a customer-supplied TFTP server. Ensure that your TFTP server is running and accessible to the FortiGate unit before you enter the command.

1. Connect to the FortiGate unit through the CLI.
2. Type the following command:

```
execute vpn certificate local export tftp <cert_name> <exp_filename> <tftp_ip>
<password>
```

where:

- <cert_name> is the name of the server certificate; typing ? displays a list of installed server certificates.
 - <exp_filename> is a name for the output file.
 - <tftp_ip> is the IP address assigned to the TFTP server host interface.
3. Move the output file from the TFTP server location to the management computer for future reference.

To import a server certificate and private key - web-based manager:

1. Go to **System > Certificates** and select **Import**.
2. In **Type**, select **PKCS12 Certificate**.
3. Select **Browse**. Browse to the location on the management computer where the exported file has been saved, select the file, and then select **Open**.
4. In the **Password** field, type the password needed to upload the exported file.
5. Select **OK**, and then select **Return**.

To import a server certificate and private key - CLI:

1. Connect to the FortiGate unit through the CLI.
2. Type the following command:

```
execute vpn certificate local import tftp <file_name> <tftp_ip_address> <file_type> <Enter
for 'cer'>|<password for 'p12'>
```

For example:

```
execute vpn certificate local import tftp FGTF-extern.p12 10.1.100.253 p12 123456
```

To import separate server certificate and private key files - web-based manager

Use the following procedure to import a server certificate and the associated private key file when the server certificate request and private key were not generated by the FortiGate unit. The two files to import must be available on the management computer.

1. Go to **System > Certificates** and select **Import**.
2. In **Type**, select **Certificate**.
3. Select the **Browse** button beside the **Certificate file** field. Browse to the location on the management computer where the certificate file has been saved, select the file, and then select **Open**.
4. Select the **Browse** button beside the **Key file** field. Browse to the location on the management computer where the key file has been saved, select the file, and then select **Open**.
5. If required, in the **Password** field, type the associated password, and then select **OK**.
6. Select **Return**.

Configuring certificate-based authentication

You can configure certificate-based authentication for FortiGate administrators, SSL VPN users, and IPsec VPN users.

In Microsoft Windows 7, you can use the certificate manager to keep track of all the different certificates on your local computer. To access certificate manager, in Windows 7 press the Windows key, enter “certmgr.msc” at the search prompt, and select the displayed match. Remember that in addition to these system certificates, many applications require you to register certificates with them directly.

To see FortiClient certificates, open the FortiClient Console, and select VPN. The VPN menu has options for My Certificates (local or client) and CA Certificates (root or intermediary certificate authorities). Use Import on those screens to import certificate files from other sources.

Authenticating administrators with security certificates

You can install a certificate on the management computer to support strong authentication for administrators. When a personal certificate is installed on the management computer, the FortiGate unit processes the certificate after the administrator supplies a username and password.

To enable strong administrative authentication:

- Obtain a signed personal certificate for the administrator from a CA and load the signed personal certificate into the web browser on the management computer according to the browser documentation.
- Install the root certificate and the CRL from the issuing CA on the FortiGate unit (see [Installing a CA root certificate and CRL to authenticate remote clients on page 249](#)).
- Create a PKI user account for the administrator.
- Add the PKI user account to a firewall user group dedicated to PKI-authenticated administrators.
- In the administrator account configuration, select **PKI** as the account **Type** and select the **User Group** to which the administrator belongs.

Support exact match for subject and CN fields in peer user

In order to avoid any unintentional admin access by regular users, administrators can specify which way a peer user authenticates.

When searching for a matching certificate, use the commands below to control how to find matches in the certificate subject name (`subject-match`) or the cn attribute (`cn-match`) of the certificate subject name. This match can be any string (`substring`) or an exact match (`value`) of the cn attribute value.

To determine certificate subject name matches - CLI:

```
config vpn certificate setting
  edit <name>
    set subject-match {substring | value}
    set cn-match {substring | value}
  next
end
```

Authenticating SSL VPN users with security certificates

While the default self-signed certificates can be used for HTTPS connections, it is preferable to use the X.509 server certificate to avoid the redirection as it can be misinterpreted as possible session hijacking. However, the server certificate method is more complex than self-signed security certificates. Also the warning message is typically displayed for the initial connection, and future connections will not generate these messages.

X.509 certificates can be used to authenticate IPsec VPN peers or clients, or SSL VPN clients. When configured to authenticate a VPN peer or client, the FortiGate unit prompts the VPN peer or client to authenticate itself using the X.509 certificate. The certificate supplied by the VPN peer or client must be verifiable using the root CA certificate installed on the FortiGate unit in order for a VPN tunnel to be established.

To enable certificate authentication for an SSL VPN user group:

1. Install a signed server certificate on the FortiGate unit and install the corresponding root certificate (and CRL) from the issuing CA on the remote peer or client.
2. Obtain a signed group certificate from a CA and load the signed group certificate into the web browser used by each user. Follow the browser documentation to load the certificates.
3. Install the root certificate and the CRL from the issuing CA on the FortiGate unit (see [Installing a CA root certificate and CRL to authenticate remote clients on page 249](#)).
4. Create a PKI user for each SSL VPN user. For each user, specify the text string that appears in the Subject field of the user's certificate and then select the corresponding CA certificate.
5. Use the `config user peergrp` CLI command to create a peer user group. Add to this group all of the SSL VPN users who are authenticated by certificate.
6. Go to **Policy & Objects > IPv4 Policy**.
7. Edit the SSL-VPN security policy.
8. Select the user group created earlier in the **Source User(s)** field.
9. Select **OK**.

Authenticating IPsec VPN users with security certificates

To require VPN peers to authenticate by means of a certificate, the FortiGate unit must offer a certificate to authenticate itself to the peer.

To enable the FortiGate unit to authenticate itself with a certificate:

1. Install a signed server certificate on the FortiGate unit.
See [To install or import the signed server certificate - web-based manager on page 249](#).
2. Install the corresponding CA root certificate on the remote peer or client. If the remote peer is a FortiGate unit, see [To install a CA root certificate on page 249](#).
3. Install the certificate revocation list (CRL) from the issuing CA on the remote peer or client. If the remote peer is a FortiGate unit, see [To import a certificate revocation list on page 249](#).
4. In the VPN phase 1 configuration, set **Authentication Method** to **Signature** and from the **Certificate Name** list select the certificate that you installed in Step 1.

To authenticate a VPN peer using a certificate, you must install a signed server certificate on the peer. Then, on the FortiGate unit, the configuration depends on whether there is only one VPN peer or if this is a dialup VPN that can be multiple peers.

To configure certificate authentication of a single peer

1. Install the CA root certificate and CRL.
2. Create a PKI user to represent the peer. Specify the text string that appears in the Subject field of the user's certificate and then select the corresponding CA certificate.
3. In the VPN phase 1 **Peer Options**, select **peer certificate** for **Accept Types** field and select the PKI user that you created in the **Peer certificate** field.

To configure certificate authentication of multiple peers (dialup VPN)

1. Install the corresponding CA root certificate and CRL.
2. Create a PKI user for each remote VPN peer. For each user, specify the text string that appears in the Subject field of the user's certificate and then select the corresponding CA certificate.
3. Use the `config user peergrp` CLI command to create a peer user group. Add to this group all of the PKI users who will use the IPsec VPN.

In the VPN phase 1 **Peer Options**, select **peer certificate group** for **Accept Types** field and select the PKI user group that you created in the **Peer certificate group** field.

Support for per-VDOM certificates

The CA and local certificate configuration is available per-VDOM. When an admin uploads a certificate to a VDOM, it will only be accessible inside that VDOM. When an admin uploads a certificate to global, it will be accessible to all VDOMs and global.

There are factory default certificates such as Fortinet_CA_SSL, Fortinet_SSL, Fortinet_Wifi, and Fortinet_Factory. These certificates are moved to per-VDOM and automatically generated when a new VDOM is created.



The Fortinet_Firmware certificate has been removed and all the attributes that use Fortinet_Firmware now use Fortinet_Factory.

CLI changes

Two new attributes `range` and `source` have been added:

`range` can be global or per-VDOM, if the certificate file is imported from global, it is a global certificate. If the certificate file is imported from a VDOM, it is VDOM certificate.

`source` can be either `factory`, `user`, or `fortiguard`:

- `factory`: The factory certificate file with FortiOS version, this includes: Fortinet_CA_SSL, Fortinet_SSL, PositiveSSL_CA, Fortinet_Wifi, Fortinet_Factory.
- `user`: Certificate file imported by the user.
- `fortiguard`: Certificate file imported from FortiGuard.

```
config certificate local
  edit Fortinet_Factory
    set range {global | vdom}
    set source {factory | user | fortiguard}
  end
end
```


GUI changes

Global and new VDOMs have the following factory default certificates:

Name	Subject	Comments	Issuer	Expires	Status	Source	Ref.
Certificates (12)							
Fortinet_SSL_ECDSA	C = US, CN = FGT6HD3916800525, L = Sunnyvale, O = Fortinet, ST = California, emailAddress = support@fortinet.com, OU = FortiGate	This certificate is embedded in...	Fortinet	2027-03-10 18:39:25 GMT	OK	Factory	0
Fortinet_SSL_ECDSA256	C = US, CN = FGT6HD3916800525, L = Sunnyvale, O = Fortinet, ST = California, emailAddress = support@fortinet.com, OU = FortiGate	This certificate is embedded in...	Fortinet	2027-03-10 19:39:00 GMT	OK	Factory	1
Fortinet_SSL	C = US, CN = FGT6HD3916800525, L = Sunnyvale, O = Fortinet, ST = California, emailAddress = support@fortinet.com, OU = FortiGate	This certificate is embedded in...	Fortinet	2027-03-10 18:39:24 GMT	OK	Factory	2
Fortinet_SSL_DSA	C = US, CN = FGT6HD3916800525, L = Sunnyvale, O = Fortinet, ST = California, emailAddress = support@fortinet.com, OU = FortiGate	This certificate is embedded in...	Fortinet	2027-03-10 18:39:25 GMT	OK	Factory	0
Fortinet_SSL_DSA1024	C = US, CN = FGT6HD3916800525, L = Sunnyvale, O = Fortinet, ST = California, emailAddress = support@fortinet.com, OU = FortiGate	This certificate is embedded in...	Fortinet	2027-03-10 19:38:59 GMT	OK	Factory	1
Fortinet_SSL_DSA2048	C = US, CN = FGT6HD3916800525, L = Sunnyvale, O = Fortinet, ST = California, emailAddress = support@fortinet.com, OU = FortiGate	This certificate is embedded in...	Fortinet	2027-03-10 19:39:00 GMT	OK	Factory	1
Fortinet_Factory	C = US, CN = FGT6HD3916800525, L = Sunnyvale, O = Fortinet, ST = California, emailAddress = support@fortinet.com, OU = FortiGate	This certificate is embedded in...	Fortinet	2038-01-19 03:14:07 GMT	OK	Factory	2
Fortinet_SSL_ECDSA384	C = US, CN = FGT6HD3916800525, L = Sunnyvale, O = Fortinet, ST = California, emailAddress = support@fortinet.com, OU = FortiGate	This certificate is embedded in...	Fortinet	2027-03-10 19:39:00 GMT	OK	Factory	1
Fortinet_SSL_RSA	C = US, CN = FGT6HD3916800525, L = Sunnyvale, O = Fortinet, ST = California, emailAddress = support@fortinet.com, OU = FortiGate	This certificate is embedded in...	Fortinet	2027-03-10 18:39:24 GMT	OK	Factory	0
Fortinet_SSL_RSA1024	C = US, CN = FGT6HD3916800525, L = Sunnyvale, O = Fortinet, ST = California, emailAddress = support@fortinet.com, OU = FortiGate	This certificate is embedded in...	Fortinet	2027-03-10 19:38:58 GMT	OK	Factory	1
Fortinet_SSL_RSA2048	C = US, CN = FGT6HD3916800525, L = Sunnyvale, O = Fortinet, ST = California, emailAddress = support@fortinet.com, OU = FortiGate	This certificate is embedded in...	Fortinet	2027-03-10 19:38:59 GMT	OK	Factory	1
Fortinet_Wifi	C = US, CN = auth-cert.fortinet.com, L = Sunnyvale, O = Fortinet, ST = California, OU = FortiWifi	This certificate is embedded in...	Entrust, Inc.	2019-05-24 13:15:35 GMT	OK	Factory	0
Local CA Certificates (2)							
Fortinet_CA_Untrusted	C = US, CN = Fortinet Untrusted CA, L = Sunnyvale, O = Fortinet, ST = California, emailAddress = support@fortinet.com, OU = Certificate Authority	This is the default CA certific...	Fortinet	2027-03-10 18:39:23 GMT	OK	Factory	2
Fortinet_CA_SSL	C = US, CN = FGT6HD3916800525, L = Sunnyvale, O = Fortinet, ST = California, emailAddress = support@fortinet.com, OU = Certificate Authority	This is the default CA certific...	Fortinet	2027-03-10 18:39:23 GMT	OK	Factory	2
External CA Certificates (3)							
Fortinet_CA	C = US, CN = support, L = Sunnyvale, O = Fortinet, ST = California, emailAddress = support@fortinet.com, OU = Certificate Authority		Fortinet	2038-01-19 22:34:39 GMT	OK	Factory	0
Fortinet_Wifi_CA	C = US, OU = (c) 2012 Entrust, Inc. - for authorized use only, O = Entrust, Inc., CN = Entrust Certification Authority - L1K		Entrust, Inc.	2030-12-05 19:43:56 GMT	OK	Factory	0
Fortinet_Wifi_CA2	C = US, OU = (c) 2009 Entrust, Inc. - for authorized use only, O = Entrust, Inc., CN = Entrust Root Certification Authority - G2		Entrust, Inc.	2024-09-23 01:31:53 GMT	OK	Factory	0

These certificates are created automatically when a new VDOM is created, with every VDOM having its own versions of these certificates.

Example — Generate a CSR on the FortiGate unit

This example follows all the steps required to create and install a local certificate on the FortiGate unit, without using CA software.

The FortiGate unit is called **myFortiGate60**, and is located at 10.11.101.101 (a private IP address) and <http://myfortigate.example.com>. Mr. John Smith (john.smith@myfortigate.example.com) is the IT administrator for this FortiGate unit, and the unit belongs to the Sales department located in Greenwich, London, England.

To generate a certificate request on the FortiGate unit - web-based manager:

1. Go to **System > Certificates**.
2. Select **Generate**.
3. In the **Certificate Name** field, enter `myFortiGate60`.



Do not include spaces in the certificate name. This will ensure compatibility of a signed certificate as a PKCS12 file to be exported later on if required.

Since the IP address is private, we will use the FQDN instead.

4. Select **Domain Name**, and enter `http://myfortigate.example.com`.
5. Enter values in the **Optional Information** area to further identify the FortiGate unit.

Organization Unit	Sales
Organization	Example.com
Locality (City)	Greenwich

State/Province	London
Country	England
e-mail	john.smith@myfortigate.example.com

6. From the **Key Type** list, select **RSA** or **Elliptic Curve**.
7. If **RSA** is selected, set **Key Size** to **2048 Bit**. If **Elliptic Curve** is selected, set **Curve Name** to **secp256r1**.
8. In **Enrollment Method**, select **File Based** to generate the certificate request
9. Select **OK**.
The request is generated and displayed in the **Local Certificates** list with a status of **PENDING**.
10. Select the **Download** button to download the request to the management computer.
11. In the **File Download** dialog box, select **Save** and save the Certificate Signing Request on the local file system of the management computer.
12. Name the file and save it on the local file system of the management computer.

Example — Generate and Import CA certificate with private key pair on OpenSSL

This example explains how to generate a certificate using OpenSSL on MS Windows. OpenSSL is available for Linux and Mac OS as well, however their terminology will vary slightly from what is presented here.

Assumptions

Before starting this procedure, ensure that you have downloaded and installed OpenSSL on Windows. One source is: <http://www.slproweb.com/products/Win32OpenSSL.html>.

Generating and importing the CA certificate and private key

The two following procedures will generate a CA certificate file and private key file, and then import it to the FortiGate unit as a local certificate.

To generate the private key and certificate

1. At the Windows command prompt, go to the OpenSSL bin directory. If you installed to the default location this will be the command:

```
cd c:\OpenSSL-Win32\bin
```

2. Enter the following command to generate the private key. You will be prompted to enter your PEM pass phrase. Choose something easy to remember such as fortinet123.

```
openssl genrsa -aes256 -out fgtpcapriv.key 2048
```

This command generates an RSA AES256 2048-bit encryption key.

3. The following command will generate the certificate using the key from the previous step.

```
openssl req -new -x509 -days 3650 -extensions v3_ca -key fgtpcapriv.key -out fgtpca.crt
```

This step generates an X509 CA certificate good for 10 years that uses the key generated in the

previous step. The certificate filename is `fgtca.crt`.

You will be prompted to enter information such as PEM Pass Phrase from the previous step, Country Name, State, Organization Name, Organizational Unit (such as department name), Common Name (the FQDN), and Email Address.

To import the certificate to the FortiGate unit - web-based manager:

1. Go to **System > Certificates**.
2. Select **Import > Local Certificate**.
3. Select **Certificate** for **Type**.
Fields for Certificate file, Key file, and Password are displayed.
4. For **Certificate file**, enter `c:\OpenSSL-Win32\bin\fgtca.crt`.
5. For **Key file**, enter `c:\OpenSSL-Win32\bin\fgtcapriv.key`.
6. For **Password**, enter the PEM Pass Phrase you entered earlier, such as `fortinet123`.
7. Select **OK**.

The Certificate will be added to the list of Local Certificates and be ready for use. It will appear in the list as the filename you uploaded — `fgtca`. You can add comments to this certificate to make it clear where its from and how it is intended to be used. If you download the certificate from FortiOS, it is a `.CER` file.

It can now be used in [Authenticating IPsec VPN users with security certificates on page 256](#), and [Authenticating SSL VPN users with security certificates on page 256](#).

Example — Generate an SSL certificate in OpenSSL

This example explains how to generate a CA signed SSL certificate using OpenSSL on MS Windows. OpenSSL is available for Linux and Mac OS as well, however their terminology will vary slightly from what is presented here.

In this example, you will:

- Generate a CA signed SSL certificate
- Generate a self-signed SSL certificate
- Import the SSL certificate into FortiOS

Assumptions

- Before starting this procedure, ensure that you have downloaded and installed OpenSSL on MS Windows. One download source is <http://www.slproweb.com/products/Win32OpenSSL.html>.

Generating a CA signed SSL certificate

This procedure assumes that you have already completed [Example — Generate and Import CA certificate with private key pair on OpenSSL on page 259](#) successfully.

To generate the CA signed SSL certificate:

1. At the Windows command prompt, go to the OpenSSL bin directory. If you installed to the default location this will be the following command:

```
cd c:\OpenSSL-Win32\bin
```

2. Enter the following command to generate the private key. You will be prompted to enter your PEM pass phrase. Choose something easy to remember such as fortinet.

```
openssl genrsa -aes256 -out fgtssl.key 2048
```

This command generates an RSA AES256 2048-bit encryption key.

3. Create a certificate signing request for the SSL certificate. This step requires you to enter the information listed in step 3 of the previous example — [To generate the private key and certificate](#). You can leave the Challenge Password blank.

```
openssl req -new -sha256 -key fgtssl.key -out fgtssl.csr
```



Most Certificate Authorities will ignore the value that is set in the CSR and use whatever value they are set to use in their configuration. This means that the client will likely need to modify their openssl.conf file to use SHA-256 (or another SHA-2 variant).

4. Using the CSR from the previous step, you can now create the SSL certificate using the CA certificate that was created in [Example — Generate and Import CA certificate with private key pair on OpenSSL](#).

```
openssl x509 -req -days 365 -in fgtssl.csr -CA fgtca.crt -CAkey fgtpcapriv.key -set_serial 01 -out fgtssl.crt
```

This will generate an X.509 certificate good for 365 days signed by the CA certificate fgtca.crt.

Generating a self-signed SSL certificate

This procedure does not require any existing certificates.

1. At the Windows command prompt, go to the OpenSSL bin directory. If you installed to the default location this will be the following command:

```
cd c:\OpenSSL-Win32\bin
```

2. Enter the following command to generate the private key. You will be prompted to enter your PEM pass phrase. Choose something easy to remember such as fortinet.

```
openssl genrsa -aes256 -out fgtssl.key 2048
openssl req -new -key fgtssl.key -out fgtssl.csr
openssl x509 -req -days 365 -in fgtssl.csr -signkey fgtssl.key -out fgtssl.crt
```

These commands:

- generate an RSA AES256 2048-bit private key,
- generate an SSL certificate signing request, and
- sign the CSR to generate an SSL .CRT certificate file.

Import the SSL certificate into FortiOS

To import the certificate to FortiOS- web-based manager

1. Go to **System > Certificates**.
2. Select **Import > Local Certificate**.
3. Select **Certificate** for **Type**.
Fields for Certificate file, Key file, and Password are displayed.
4. For **Certificate file**, enter `c:\OpenSSL-Win32\bin\fgtssl.crt`.
5. For **Key file**, enter `c:\OpenSSL-Win32\bin\fgtssl.key`.
6. For **Password**, enter the PEM Pass Phrase you entered, such as `fortinet`.
7. Select **OK**.

The SSL certificate you just uploaded can be found under **System > Certificates** under the name of the file you uploaded — `fgtssl`.

To confirm the certificate is uploaded properly - CLI:

```
config vpn certificate local
  edit fgtssl
    get
  end
```

The `get` command will display all the certificate's information. If it is not there or the information is not correct, you will need to remove the corrupted certificate (if it is there) and upload it again from your PC.

To use the new SSL certificate - CLI

```
config vpn ssl settings
  set servercert fgtssl
end
```

This assigns the `fgtssl` certificate as the SSL server certificate. For more information see the FortiOS Handbook SSL VPN guide.

Single sign-on using a FortiAuthenticator unit

If you use a FortiAuthenticator unit in your network as a single sign-on agent,

- Users can authenticate through a web portal on the FortiAuthenticator unit.
- Users with FortiClient Endpoint Security installed can be automatically authenticated by the FortiAuthenticator unit through the FortiClient SSO Mobility Agent.

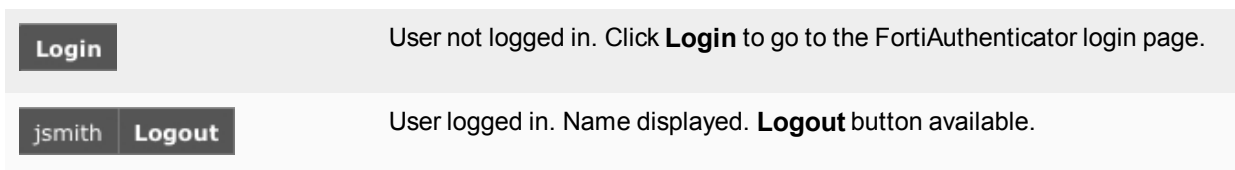
The FortiAuthenticator unit can integrate with external network authentication systems such as RADIUS and LDAP to gather user logon information and send it to the FortiGate unit.

User's view of FortiAuthenticator SSO authentication

There are two different ways users can authenticate through a FortiAuthenticator unit.

Users without FortiClient Endpoint Security - SSO widget

To log onto the network, the user accesses the organization's web page with a web browser. Embedded on that page is a simple logon widget, like this:



The SSO widget sets a cookie on the user's browser. When the user browses to a page containing the login widget, the FortiAuthenticator unit recognizes the user and updates its database if the user's IP address has changed. The user will not need to re-authenticate until the login timeout expires, which can be up to 30 days.

Users with FortiClient Endpoint Security - FortiClient SSO Mobility Agent

The user simply accesses resources and all authentication is performed transparently with no request for credentials. IP address changes, such as those due to WiFi roaming, are automatically sent to the FortiAuthenticator unit. When the user logs off or otherwise disconnects from the network, the FortiAuthenticator unit is aware of this and deauthenticates the user.

The FortiClient SSO Mobility Agent, a feature of FortiClient Endpoint Security v5.0, must be configured to communicate with the appropriate FortiAuthenticator unit. After that, the agent automatically provides user name and IP address information to the FortiAuthenticator unit for transparent authentication.

Administrator's view of FortiAuthenticator SSO authentication

You can configure either or both of these authentication types on your network.

SSO widget

You need to configure the Single Sign-On portal on the FortiAuthenticator unit. Go to **Fortinet SSO Methods > SSO > Portal Services** to do this. Copy the **Embeddable login widget** code for use on your organization's

home page. Identity-based security policies on the FortiGate unit determine which users or groups of users can access which network resources.

FortiClient SSO Mobility Agent

Your users must be running at least FortiClient Endpoint Security v5.0 to make use of this type of authentication.

On the FortiAuthenticator unit, you need to select **Enable FortiClient SSO Mobility Agent Service**, optionally select **Enable Authentication** and choose a **Secret key**. Go to **Fortinet SSO Methods > SSO > General**. You need to provide your users the FortiAuthenticator IP address and secret key so that they can configure the FortiClient SSO Mobility Agent on their computers. See [Configuring the FortiGate unit on page 265](#).

Configuring the FortiAuthenticator unit

The FortiAuthenticator unit can poll FortiGate units, Windows Active Directory, RADIUS servers, LDAP servers, and FortiClients for information about user logon activity.

To configure FortiAuthenticator polling:

1. Go to **Fortinet SSO Methods > SSO > General**.
2. In the **FortiGate** section, leave the Listening port at 8000, unless your network requires you to change this. The FortiGate unit must allow traffic on this port to pass through the firewall. Optionally, you can set the Login Expiry time. This is the length of time users can remain logged in before the system logs them off automatically. The default is 480 minutes (8 hours).
3. Select **Enable Authentication** and enter the **Secret key**. Be sure to use the same secret key when configuring the FSSO Agent on FortiGate units.
4. In the **Fortinet Single Sign-On (FSSO)** section, enter

Enable Windows Active Directory domain controllers	Select for integration with Windows Active Directory.
Enable Radius accounting SSO clients	Select if you want to use a Remote Radius server.
Enable Syslog SSO	Select for integration with Syslog server.
Enable FortiClient SSO Mobility Agent service	Select both options to enable single sign-on by clients running FortiClient Endpoint Security. Enter the Secret key . Be sure to use the same secret key in the FortiClient Single Sign-On Mobility Agent settings.
Enable Authentication	

5. Select **OK**.

For more information, see the FortiAuthenticator Administration Guide.

Configuring the FortiGate unit

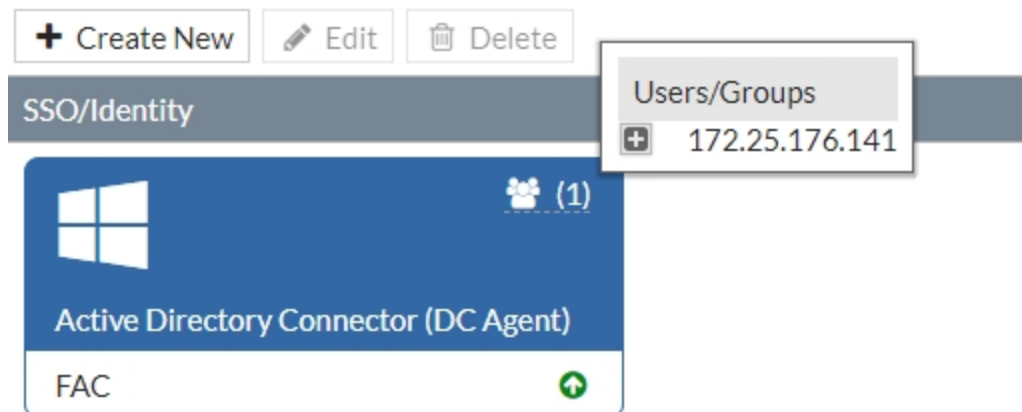
Adding a FortiAuthenticator unit as an SSO agent

On the FortiGate unit, you need to add the FortiAuthenticator unit as a Single Sign-On agent that provides user logon information.

To add a FortiAuthenticator unit as an SSO agent:

1. Go to **Security Fabric > Fabric Connectors** and select **Create New**.
2. Under **SSO/Identity**, select **Fortinet Single Sign-On Agent**.
3. Enter a **Name** for the FortiAuthenticator unit (in the example, **FAC**).
4. In **Primary FSSO Agent**, enter the IP address of the FortiAuthenticator unit and password.
On the FortiAuthenticator unit, go to **Fortinet SSO Methods > SSO > General** to define the secret key. Select **Enable Authentication**.
5. Keep **Collector Agent AD access mode** set to **Standard**, and select **OK**.

The entry is shown in the **SSO/Identity** server list, with a green arrow indicating a successful connection. Select the plus-symbol to view the list of user groups that the FortiGate has received from the FortiAuthenticator.



When you open the server, you can see the list of groups. You can use the groups in identity-based security policies.

Configuring an FSSO user group

You cannot use FortiAuthenticator SSO user groups directly in a security policy. Create an FSSO user group and add FortiAuthenticator SSO user groups to it. FortiGate FSSO user groups are available for selection in identity-based security policies.

To create an FSSO user group:

1. Go to **User & Device > User Groups** and select **Create New**.
2. Enter a **Name** for the group.
3. In Type, select **Fortinet Single Sign-On (FSSO)**.
4. Add **Members**.
The groups available to add as members are SSO groups provided by SSO agents.
5. Select **OK**.

Configuring security policies

You can create identity-based policies based on FSSO groups as you do for local user groups. For more information about security policies see the Firewall chapter.

Configuring the FortiClient SSO Mobility Agent

The user's device must have at least FortiClient Endpoint Security v5.0 installed. Only two pieces of information are required to set up the SSO Mobility Agent feature: the FortiAuthenticator unit IP address and the pre-shared secret.

The user needs to know the FortiAuthenticator IP address and pre-shared secret to set up the SSO Mobility Agent. Or, you could preconfigure FortiClient.

To configure FortiClient SSO Mobility Agent:

1. In FortiClient Endpoint Security, go to **File > Settings**.
You must run the FortiClient application as an administrator to access these settings.
2. Select **Enable single sign-on mobility agent**. Enter the FortiAuthenticator unit IP address, including the listening port number specified on the FortiAuthenticator unit.
Example: 192.168.0.99:8001. You can omit the port number if it is 8005.
3. Enter the pre-shared key.
4. Select **OK**.

Viewing SSO authentication events on the FortiGate unit

User authentication events are logged in the FortiGate event log.

Go to **Log & Report > System Events**.

FortiWiFi 60CX-ADSL-A

Help Wizard Logout FORTINET

Refresh Download Raw Log Column Settings Log location: Disk

System
Policy
Firewall Objects
UTM Security Profiles
VPN
User
WAN Opt. & Cache
WiFi Controller
Log&Report

Traffic Log
Forward Traffic
Local Traffic
Multicast Traffic
Invalid Packets
Event Log
System
Router
VPN
User
WAN Opt. & Cache
WiFi

#	Date/Time	Level	User	Action	Message	Group
1	1 minute ago	notice	KADIMUNDI	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user KADIMUNDI logged on 192.168.1.100"	
2	1 minute ago	notice	aventress	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user aventress logged on 192.168.0.150"	
3	2 minutes ago	notice	KADIMUNDI	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user KADIMUNDI logged on 192.168.1.100"	
4	2 minutes ago	notice	DOMAINADMIN	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user DOMAINADMIN logged on 192.168.1.100"	
5	3 minutes ago	notice	ATANO	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user ATANO logged on 192.168.1.101"	
6	6 minutes ago	notice	DOMAINADMIN	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user DOMAINADMIN logged on 192.168.1.100"	
7	11 minutes ago	notice	DOMAINADMIN	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user DOMAINADMIN logged on 192.168.1.100"	
8	16 minutes ago	notice	DOMAINADMIN	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user DOMAINADMIN logged on 192.168.1.100"	
9	20 minutes ago	notice	DOMAINADMIN	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user DOMAINADMIN logged on 192.168.1.100"	
10	21 minutes ago	notice	DOMAINADMIN	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user DOMAINADMIN logged on 192.168.1.100"	
11	26 minutes ago	notice	DOMAINADMIN	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user DOMAINADMIN logged on 192.168.1.100"	
12	31 minutes ago	notice	DOMAINADMIN	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user DOMAINADMIN logged on 192.168.1.100"	
13	36 minutes ago	notice	DOMAINADMIN	FSSO-logout	FSSO-logout event from FortiAuthenticator_2.0: user DOMAINADMIN logged on 192.168.1.100"	

1 / 5 [Total: 248]

Date/Time	1 minute ago (Mon Oct 1 17:59:08 2012)	Level	notice
Sub Type	user	User	KADIMUNDI
Action	FSSO-logout	Message	FSSO-logout event from FortiAuthenticator_2.0: user KADIMUNDI logged on 192.168.1.100"
Src	192.168.1.100	Dst	FortiAuthenticator_2.0

Single sign-on to Windows AD

The FortiGate unit can authenticate users transparently and allow them network access based on their privileges in Windows AD. This means that users who have logged on to the network are not asked again for their credentials to access network resources through the FortiGate unit, hence the term “Single Sign-On” (SSO).

The following topics are included:

- [Introduction to SSO with Windows AD](#)
- [Configuring SSO to Windows AD](#)
- [FortiOS FSSO log messages](#)
- [Testing FSSO](#)
- [Troubleshooting FSSO](#)

Introduction to SSO with Windows AD

SSO support provided by FortiGate polling of domain controllers is simpler than the earlier method that relies on agent software installed on Windows AD network servers. No Fortinet software needs to be installed on the Windows network. The FortiGate unit needs access only to the Windows AD global catalog and event log.

When a Windows AD user logs on at a workstation in a monitored domain, the FortiGate unit:

- detects the logon event in the domain controller’s event log and records the workstation name, domain, and user,
- resolves the workstation name to an IP address,
- uses the domain controller’s LDAP server to determine which groups the user belongs to,
- and creates one or more log entries on the FortiGate unit for this logon event as appropriate.

When the user tries to access network resources, the FortiGate unit selects the appropriate security policy for the destination. The selection consists of matching the FSSO group or groups the user belongs to with the security policy or policies that match that group. If the user belongs to one of the permitted user groups associated with that policy, the connection is allowed. Otherwise the connection is denied.

Configuring SSO to Windows AD

On the FortiGate unit, security policies control access to network resources based on user groups. With Fortinet SSO, this is also true but each FortiGate user group is associated with one or more Windows AD user groups. This is how Windows AD user groups get authenticated in the FortiGate security policy.

Fortinet SSO (FSSO) sends information about Windows user logons to FortiGate units. If there are many users on your Windows AD domains, the large amount of information might affect the performance of the FortiGate unit.

To configure your FortiGate unit to operate with either a Windows AD or a Novell eDirectory FSSO install, you

- Configure LDAP access to the Windows AD global catalog. See [Configuring LDAP server access on page 269](#).
- Configure the LDAP Server as a Single Sign-On server. See [Configuring the LDAP server as an SSO server on page 270](#).
- Add Active Directory user groups to FortiGate FSSO user groups. See [Creating FSSO user groups on page 271](#).
- Create security policies for FSSO-authenticated groups. See [Creating security policies on page 271](#).

- Optionally, specify a guest protection profile to allow guest access. See [Enabling guest access through FSSO security policies on page 272](#)

Configuring LDAP server access

The FortiGate unit needs access to the domain controller's LDAP server to retrieve user group information.

The LDAP configuration on the FortiGate unit not only provides access to the LDAP server, it sets up the retrieval of Windows AD user groups for you to select in FSSO. The LDAP Server configuration, found under **User & Device > LDAP Servers**, includes a function to preview the LDAP server's response to your distinguished name query. If you already know the appropriate Distinguished Name (DN) and User DN settings, you may be able to skip some of the following steps.

To add an LDAP server - web-based manager:

1. Go to **User & Device > LDAP Servers** and select **Create New**.
2. Enter the **Server IP/Name** and **Server Port** (default 389).
3. In the **Common Name Identifier** field, enter **sAMAccountName**. The default common name identifier is **cn**. This is correct for most LDAP servers. However some servers use other identifiers such as **uid**.
4. In the **Distinguished Name** field, enter your organization distinguished name. In this example, Distinguished Name is **dc=techdoc,dc=local**
5. Select **Fetch DN**, this will fetch the Windows AD directory.

The screenshot shows the 'LDAP Servers' configuration page in the FortiGate web-based manager. The following fields are populated:

- Name:** LDAP
- Server IP/Name:** 10.10.20.3
- Server Port:** 389
- Common Name Identifier:** sAMAccountName
- Distinguished Name:** dc=techdoc,dc=local
- Bind Type:** Regular (selected)
- User DN:** (empty)
- Password:** (empty)
- Secure Connection:** ☐
- Test:** (button)

An expanded window titled 'LDAP Distinguished Name Query' displays an 'LDAP Tree' with the following structure:

- dc=techdoc,dc=local (root)
 - CN=Computers
 - CN=ForeignSecurityPrincipals
 - CN=Managed Service Accounts
 - CN=Program Data
 - CN=System
 - CN=Users
 - OU=Domain Controllers

6. Set **Bind Type** to **Regular**.
7. In the **User DN** field, enter the administrative account name that you created for FSSO. For example, if the account is administrator, enter "administrator@techdoc.local".
8. Enter the administrative account password in the **Password** field.
9. Optionally select **Secure Connection**.
 - In the **Protocol** field, select **LDAPS** or **STARTTLS**.
 - In the **Certificate** field, select the appropriate certificate for authentication.

Note that you need to configure the Windows AD for secure connection accordingly.

10. Select **OK**.

11. Test your configuration by selecting the **Test** button. A successful message confirming the right settings appears.

The screenshot shows a configuration window for LDAP. At the top, a blue banner reads "Successful". Below it, the following fields are visible:

- Name:** LDAP
- Server IP/Name:** 10.10.20.3
- Server Port:** 389
- Common Name Identifier:** sAMAccountName
- Distinguished Name:** dc=techdoc,dc=local
- Fetch DN:** A button labeled "Fetch DN".
- Bind Type:** Radio buttons for Simple, Anonymous, and Regular (Regular is selected).
- User DN:** administrator@techdoc.local
- Password:** A masked password field with dots.
- Secure Connection:** An unchecked checkbox.
- Test:** A button labeled "Test".

To configure LDAP for FSSO - CLI example:

```
config user ldap
  edit LDAP
    set server 10.10.20.3
    set cnid sAMAccountName
    set dn dc=techdoc,dc=local
    set type regular
    set username administrator@techdoc.local
    set password <your_password>
  next
end
```

Configuring the LDAP server as an SSO server

The LDAP server must be added to the FortiGate SSO configuration.

To add the LDAP server as an SSO server:

1. Go to **Security Fabric > Fabric Connectors** and select **Create New**.
2. Enter the following:

SSO/Identity	Select Poll Active Directory Server .
Server IP/Name	Server Name or IP address of the Domain Controller.
User	A Domain user name.
Password	The user's password.
LDAP Server	Select the LDAP server you added earlier.
Enable Polling	Enable.

3. Select **OK**.

Creating FSSO user groups

You cannot use Windows or Novell groups directly in FortiGate security policies. You must create FortiGate user groups of the FSSO type and add Windows or Novell groups to them.

To create a user group for FSSO authentication - web-based manager:

1. Go to **User & Device > User Groups** and select **Create New**.
The **New User Group** dialog box opens.
2. In the **Name** box, enter a name for the group, `FSSO_Internet_users` for example.
3. In **Type**, select **Fortinet Single Sign-On (FSSO)**.
4. In **Members**, select the required **FSSO** groups.
5. Select **OK**.

To create the `FSSO_Internet-users` user group - CLI

```
config user group
  edit FSSO_Internet_users
    set group-type fsso-service
    set member CN=Engineering,cn=users,dc=office,dc=example,dc=com
      CN=Sales,cn=users,dc=office,dc=example,dc=com
  end
```

Default FSSO group

`SSO_Guest_users` is a default user group enabled when FSSO is configured. It allows guest users on the network who do not have an FSSO account to authenticate and have access to network resources. See [Enabling guest access through FSSO security policies on page 272](#).

Creating security policies

Policies that require FSSO authentication are very similar to other security policies. Using identity-based policies, you can configure access that depends on the FSSO user group. This allows each FSSO user group to have its own level of access to its own group of services

In this situation, Example.com is a company that has its employees and authentication servers on an internal network. The FortiGate unit intercepts all traffic leaving the internal network and requires FSSO authentication to access network resources on the Internet. The following procedure configures the security policy for FSSO authentication. FSSO is installed and configured including the RADIUS server, FSSO Collector agent, and user groups on the FortiGate

For the following procedure, the internal interface is `port1` and the external interface connected to the Internet is `port2`. There is an address group for the internal network called `company_network`. The FSSO user group is called `fsso_group`, and the FSSO RADIUS server is `fsso_rad_server`.

To configure an FSSO authentication security policy - web-based manager:

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information.

Incoming Interface	port1
Source Address	company_network
Source User(s)	fsso_group
Outgoing Interface	port2
Destination Address	all
Schedule	always
Service	HTTP, HTTPS, FTP, and Telnet
Action	ACCEPT
NAT	ON
UTM Security Profiles	ON for AntiVirus, IPS, Web Filter, and Email Filter, all using default profiles.
Log Allowed Traffic	ON. Select Security Events .

3. Select **OK**.
4. Ensure the FSSO authentication policy is higher in the policy list than more general policies for the same interfaces.

To create a security policy for FSSO authentication - CLI:

```
config firewall policy
  edit 0
    set srcintf port1
    set dstintf port2
    set srcaddr company_network
    set dstaddr all
    set action accept
    set groups fsso_group
    set schedule always
    set service HTTP HTTPS FTP TELNET
    set nat enable
  end
```

Here is an example of how this FSSO authentication policy is used. Example.com employee on the internal company network logs on to the internal network using their RADIUS username and password. When that user attempts to access the Internet, which requires FSSO authentication, the FortiGate authentication security policy intercepts the session, checks with the FSSO Collector agent to verify the user's identity and credentials, and then if everything is verified the user is allowed access to the Internet.

Enabling guest access through FSSO security policies

You can enable guest users to access FSSO security policies. Guests are users who are unknown to Windows AD and servers that do not logon to a Windows AD domain.

To enable guest access in your FSSO security policy, add an identity-based policy assigned to the built-in user group `SSO_Guest_Users`. Specify the services, schedule and UTM profiles that apply to guest users — typically guests have access to a reduced set of services. See [Creating security policies on page 271](#).

FortiOS FSSO log messages

There are two types of FortiOS log messages — firewall and event. FSSO related log messages are generated from authentication events. These include user logon and log off events, and NTLM authentication events. These log messages are central to network accounting policies, and can also be useful in troubleshooting issues. For more information on firewall logging, see [Enabling security logging on page 216](#). For more information on logging, see the FortiOS Handbook Logging and Reporting guide.

Enabling authentication event logging

For the FortiGate unit to log events, that specific type of event must be enabled under logging.

When VDOMs are enabled certain options may not be available, such as CPU and memory usage events. You can enable event logs only when you are logged on to a VDOM; you cannot enable event logs globally.

To ensure you log all the events needed, set the minimum log level to Notification or Information. Firewall logging requires Notification as a minimum. The closer to Debug level, the more information will be logged.

To enable event logging:

1. Go to **Log & Report > Log Settings**.
2. Under **Log Settings**, set **Event Logging** to **Customize** and select

System activity event	All system-related events, such as ping server failure and gateway status.
User activity event	All administration events, such as user logins, resets, and configuration updates.

3. Select **Apply**.

List of FSSO related log messages

Message ID	Severity	Description
43008	Notification	Authentication was successful
43009	Notification	Authentication session failed
43010	Warning	Authentication locked out
43011	Notification	Authentication timed out
43012	Notification	FSSO authentication was successful
43013	Notification	FSSO authentication failed
43014	Notification	FSSO user logged on
43015	Notification	FSSO user logged off

Message ID	Severity	Description
43016	Notification	NTLM authentication was successful
43017	Notification	NTLM authentication failed

For more information on logging, see the FortiOS Handbook Logging and Reporting guide.

Extra filter options for security events

Logon events are detected by the FSSO CA by monitoring the Security Event logs. Additional logon event filters, such as ServiceName and ServiceID, have been implemented so as to avoid instances of conflicting security events, where existing user information could be overwritten.

Testing FSSO

Once FSSO is configured, you can easily test to ensure your configuration is working as expected. For additional FSSO testing, see [Troubleshooting FSSO on page 274](#).

1. Logon to one of the stations on the FSSO domain, and access an Internet resource.
2. Connect to the CLI of the FortiGate unit, and if possible log the output.
3. Enter the following command: `diagnose debug authd fssolist`
4. Check the output. If FSSO is functioning properly you will see something similar to the following:

```
----FSSO logons----
IP: 192.168.1.230 User: ADMINISTRATOR Groups: VLAD-AD/DOMAIN USERS
IP: 192.168.1.240 User: ADMINISTRATOR Groups: VLAD-AD/DOMAIN USERS
Total number of users logged on: 2
----end of FSSO logons----
```

The exact information will vary based on your installation.

5. Check the FortiGate event log, for FSSO-auth action or other FSSO related events with FSSO information in the message field.
6. To check server connectivity, run the following commands from the CLI:

```
FGT# diagnose debug enable
FGT# diagnose debug authd fssoserver-status
FGT# Server Name Connection Status
-----
SBS-2003 connected
```

Troubleshooting FSSO

When installing, configuring, and working with FSSO some problems are quite common. A selection of these problems follows including explanations and solutions.

Some common Windows AD problems include:

- [General troubleshooting tips for FSSO](#)
- [Users on a particular computer \(IP address\) can not access the network](#)
- [Guest users do not have access to network](#)

General troubleshooting tips for FSSO

The following tips are useful in many FSSO troubleshooting situations.

- Ensure all firewalls are allowing the FSSO required ports through.
FSSO has a number of required ports that must be allowed through all firewalls or connections will fail. These include: ports 139, 389 (LDAP), 445, 636 (LDAP).
- Ensure there is at least 64kbps bandwidth between the FortiGate unit and domain controllers. If there is insufficient bandwidth, some FSSO information might not reach the FortiGate unit. The best solution is to configure traffic shaping between the FortiGate unit and the domain controllers to ensure that the minimum bandwidth is always available.

Users on a particular computer (IP address) can not access the network

Windows AD Domain Controller agent gets the username and workstation where the logon attempt is coming from. If there are two computers with the same IP address and the same user trying to logon, it is possible for the authentication system to become confused and believe that the user on computer_1 is actually trying to access computer_2.

Windows AD does not track when a user logs out. It is possible that a user logs out on one computer, and immediately logs onto a second computer while the system still believes the user is logged on the original computer. While this is allowed, information that is intended for the session on one computer may mistakenly end up going to the other computer instead. The result would look similar to a hijacked session.

Solutions

- Ensure each computer has separate IP addresses.
- Encourage users to logout on one machine before logging onto another machine.
- If multiple users have the same username, change the usernames to be unique.
- Shorten timeout timer to flush inactive sessions after a shorter time.

Guest users do not have access to network

A group of guest users was created, but they don't have access.

Solution

The group of the guest users was not included in a policy, so they do not fall under the guest account. To give them access, associate their group with a security policy.

Additionally, there is a default group called `SSO_Guest_Users`. Ensure that group is part of an identity-based security policy to allow traffic.

Agent-based FSSO

FortiOS can provide single sign-on capabilities to Windows AD, Citrix, VMware Horizon, Novell eDirectory, or, as of FortiOS 5.4, Microsoft Exchange users with the help of agent software installed on these networks. The agent software sends information about user logons to the FortiGate unit. With user information such as IP address and user group memberships from the network, FortiGate security policies can allow authenticated network access to users who belong to the appropriate user groups without requesting their credentials again.

For Windows AD networks, FortiGate units can provide SSO capability without agent software by directly polling the Windows AD domain controllers. For information about this type of SSO, see [Single sign-on to Windows AD on page 268](#).

The following topics are included:

- [Introduction to agent-based FSSO](#)
- [FSSO NTLM authentication support](#)
- [Agent installation](#)
- [Configuring the FSSO collector agent for Windows AD](#)
- [Configuring the FSSO TS agent for Citrix](#)
- [Configuring FSSO with Novell networks](#)
- [Configuring FSSO advanced settings](#)
- [Configuring FSSO on FortiGate units](#)
- [FortiOS FSSO log messages](#)
- [Testing FSSO](#)
- [Troubleshooting FSSO](#)

Introduction to agent-based FSSO

Fortinet Single Sign-On (FSSO), through agents installed on the network, monitors user logons and passes that information to the FortiGate unit. When a user logs on at a workstation in a monitored domain, FSSO:

- detects the logon event and records the workstation name, domain, and user,
- resolves the workstation name to an IP address,
- determines which user groups the user belongs to,
- sends the user logon information, including IP address and groups list, to the FortiGate unit,
- and creates one or more log entries on the FortiGate unit for this logon event as appropriate.

When the user tries to access network resources, the FortiGate unit selects the appropriate security policy for the destination. If the user belongs to one of the permitted user groups associated with that policy, the connection is allowed. Otherwise the connection is denied.



FSSO can also provide NTLM authentication service for requests coming from FortiGate. SSO is very convenient for users, but may not be supported across all platforms. NTLM is not as convenient, but it enjoys wider support. See [FSSO NTLM authentication support on page 282](#).

Introduction to FSSO agents

There are several different FSSO agents that can be used in an FSSO implementation:

- Domain Controller (DC) agent
- eDirectory agent
- Citrix/Terminal Server (TS) agent
- Collector agent

Consult the latest FortiOS and FSSO Release Notes for operating system compatibility information.

Domain Controller (DC) agent

The Domain Controller (DC) agent must be installed on every domain controller if you will use DC Agent mode, but is not required if you use Polling mode. See [FSSO for Windows AD on page 278](#).

eDirectory agent

The eDirectory agent is installed on a Novell network to monitor user logons and send the required information to the FortiGate unit. It functions much like the Collector agent on a Windows AD domain controller. The agent can obtain information from the Novell eDirectory using either the Novell API or LDAP.

Terminal Server (TS) agent

Terminal Server (TS) agent can be installed on a Citrix or VMware Horizon 7.4 terminal server to monitor user logons in real time. It functions much like the DC Agent on a Windows AD domain controller.

Collector agent

This agent is installed as a service on a server in the Windows AD network to monitor user logons and send the required information to the FortiGate unit. The Collector agent can collect information from

- Domain Controller agent (Windows AD)
- TS agent (Citrix or VMware Horizon Terminal Server)

In a Windows AD network, the Collector agent can optionally obtain logon information by polling the AD domain controllers. In this case, DC agents are not needed.

The Collector can obtain user group information from the DC agent or optionally, a FortiGate unit can obtain group information directly from AD using Lightweight Directory Access Protocol (LDAP).

On a Windows AD network, the FSSO software can also serve NT LAN Manager (NTLM) requests coming from client browsers (forwarded by the FortiGate unit) with only one or more Collector agents installed. See [FSSO NTLM authentication support on page 282](#).

The CA is responsible for DNS lookups, group verification, workstation checks, and as mentioned FortiGate updates of logon records. The FSSO Collector Agent sends Domain Local Security Group and Global Security Group information to FortiGate units. The CA communicates with the FortiGate over TCP port 8000 and it listens on UDP port 8002 for updates from the DC agents.

The FortiGate unit can have up to five CAs configured for redundancy. If the first on the list is unreachable, the next is attempted, and so on down the list until one is contacted. See [Configuring FSSO on FortiGate units on page 310](#).

All DC agents must point to the correct Collector agent port number and IP address on domains with multiple DCs.



A FortiAuthenticator unit can act much like a Collector agent, collecting Windows AD user logon information and sending it to the FortiGate unit. It is particularly useful in large installations with several FortiGate units. For more information, see the [FortiAuthenticator Administration Guide](#).

FSSO for Microsoft Exchange Server

As of FortiOS 5.4, FSSO supports monitoring Microsoft Exchange Server. This is useful for situations when the user accesses the domain account to view their email, even when the client device might not be in the domain.

Support for the Exchange server is configured on the Back-end FSSO collector agent. For more information on the collector agent, see [Collector agent installation](#):

1. On the FSSO collector agent, go to **Advanced Settings > Exchange Server**.
2. Select **Add** and enter the following information and select **OK**:

Domain Name	Enter your domain name.
Server IP/Hostname	Enter the IP address or the hostname of your exchange server.
Polling forwarded event log	This option for scenarios when you do not want that CA polls the Exchange Server logs directly. In this case you need to configure event log forwarding on the Exchange server. Exchange event logs can be forwarded to any member server. If you enable this, instead of the IP of the Exchange server configured in the previous step, you must then configure the IP of this member server. CA will then contact the member server.
Ignore Name	<p>Because CA will also check Windows log files for logon events and when a user authenticates to Exchange Server there is also a logon event in Windows event log, which CA will read and this will overwrite the Exchange Server logon event (ESEventLog) on CA. So it is recommended to set the ignore list to the domain the user belongs to.</p> <p>To do so, enter the domain name in the Ignore Name field and select Add.</p>

FSSO for Windows AD

FSSO for Windows AD requires at least one Collector agent. Domain Controller agents may also be required depending on the Collector agent working mode. There are two working modes to monitor user logon activity: DC Agent mode or Polling mode.

Collector agent DC Agent mode versus Polling mode

	DC Agent mode	Polling Mode
Installation	Complex — Multiple installations: one agent per DC plus Collector agent, requires a reboot	Easy — Only Collector agent installation, no reboot required

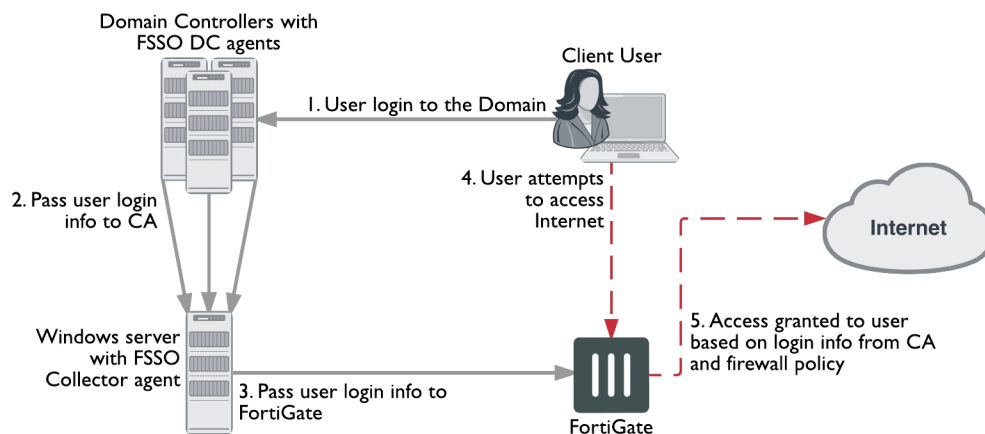
	DC Agent mode	Polling Mode
Resources	Shares resources with DC system	Has own resources
Network load	Each DC agent requires minimum 64kpbs bandwidth, adding to network load	Increase polling period during busy period to reduce network load
Level of Confidence	Captures all logons	Potential to miss a login if polling period is too great

DC Agent mode

DC Agent mode is the standard mode for FSSO. In DC Agent mode, a Fortinet authentication agent is installed on each domain controller. These DC agents monitor user logon events and pass the information to the Collector agent, which stores the information and sends it to the FortiGate unit.

The DC agent installed on the domain controllers is not a service like the Collector agent — it is a DLL file called `dcagent.dll` and is installed in the `Windows\system32` directory. It must be installed on all domain controllers of the domains that are being monitored.

FSSO in DC agent mode



DC Agent mode provides reliable user logon information, however you must install a DC agent on every domain controller. A reboot is needed after the agent is installed. Each installation requires some maintenance as well. For these reasons it may not be possible to use the DC Agent mode.

Each domain controller connection needs a minimum guaranteed 64kpbs bandwidth to ensure proper FSSO functionality. You can optionally configure traffic shapers on the FortiGate unit to ensure this minimum bandwidth is guaranteed for the domain controller connections.

Polling mode

In Polling mode there are three options — NetAPI polling, Event log polling, and Event log using WMI. All share the advantages of being transparent and agentless.

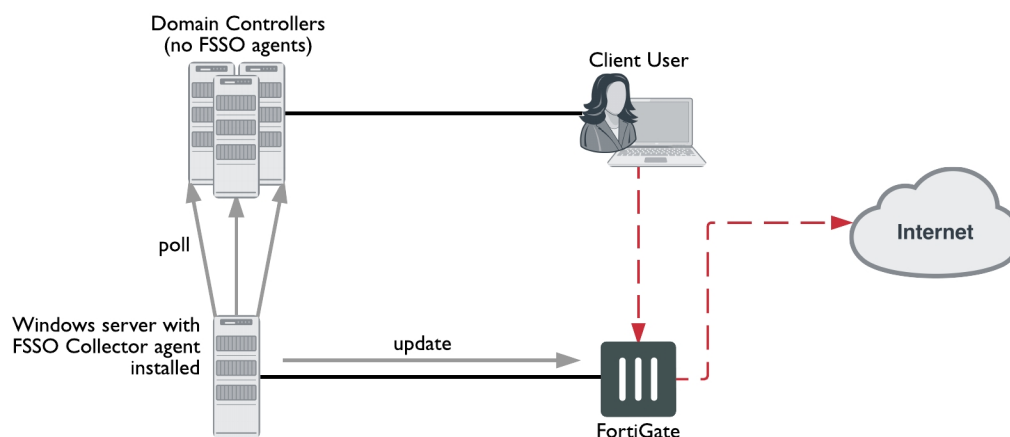
NetAPI polling is used to retrieve server logon sessions. This includes the logon event information for the Controller agent. NetAPI runs faster than Event log polling but it may miss some user logon events under heavy system load. It requires a query round trip time of less than 10 seconds.

Event log polling may run a bit slower, but will not miss events, even when the installation site has many users that require authentication. It does not have the 10 second limit on NetAPI polling. Event log polling requires fast network links. Event log polling is required if there are Mac OS users logging into Windows AD.

Event log using WMI polling: WMI is a Windows API to get system information from a Windows server, CA is a WMI client and sends WMI queries for user logon events to DC, which in this case is a WMI server. Main advantage in this mode is that CA does not need to search security event logs on DC for user logon events, instead, DC returns all requested logon events via WMI. This also reduces network load between CA and DC.

In Polling mode, the Collector agent polls port 445 of each domain controller for user logon information every few seconds and forwards it to the FortiGate unit. There are no DC Agents installed, so the Collector agent polls the domain controllers directly.

FSSO in polling mode



A major benefit of Polling mode is that no FSSO DC Agents are required. If it is not possible to install FSSO DC Agents on your domain controllers, this is the alternate configuration available to you. Polling mode results in a less complex install, and reduces ongoing maintenance. The minimum permissions required in Polling mode are to read the event log or call NetAPI.

Collector agent AD access mode - standard versus advanced

The Collector agent has two ways to access Active Directory user information. The main difference between standard and advanced mode is the naming convention used when referring to username information.

Standard mode uses regular Windows convention: Domain\Username. Advanced mode uses LDAP convention: CN=User, OU=Name, DC=Domain.

If there is no special requirement to use LDAP— best practices suggest you set up FSSO in standard mode. This mode is easier to set up, and is usually easier to maintain and troubleshoot.


Standard and advanced modes have the same level of functionality with the following exceptions:

- Users have to create Group filters on the Collector agent. This differs from advanced mode where Group filters are configured from the FortiGate unit. Fortinet strongly encourages users to create filters from CA.
- Advanced mode supports nested or inherited groups. This means that users may be a member of multiple monitored groups. Standard mode does not support nested groups so a user must be a direct member of the group being monitored.

FSSO for Citrix

Citrix users can enjoy a similar Single Sign-On experience as Windows AD users. The FSSO TS agent installed on each Citrix server provides user logon information to the FSSO Collector agent on the network. The FortiGate unit uses this information to authenticate the user in security policies.

Citrix SSO topology



Citrix users do not have unique IP addresses. When a Citrix user logs on, the TS agent assigns that user a range of ports. By default each user has a range of 200 ports.

FSSO for Novell eDirectory

FSSO in a Novell eDirectory environment works similar to the FSSO Polling mode in the Windows AD environment. The eDirectory agent polls the eDirectory servers for user logon information and forwards the information to the FortiGate unit. There is no need for the Collector agent.

When a user logs on at a workstation, FSSO:

- detects the logon event by polling the eDirectory server and records the IP address and user ID,
- looks up in the eDirectory which groups this user belongs to,
- sends the IP address and user groups information to the FortiGate unit.

When the user tries to access network resources, the FortiGate unit selects the appropriate security policy for the destination. If the user belongs to one of the permitted user groups, the connection is allowed.

FSSO is supported on the Novell E-Directory 8.8 operating system.

For a Novell network, there is only one FSSO component to install — the eDirectory agent. In some cases, you also need to install the Novell Client.

FSSO security issues

When the different components of FSSO are communicating there are some inherent security features.

FSSO installation requires an account with network admin privileges. The security inherent in these types of accounts helps ensure access to FSSO configurations is not tampered with.

User passwords are never sent between FSSO components. The information that is sent is information to identify a user including the username, group or groups, and IP address.

NTLM uses base-64 encoded packets, and uses a unique randomly generated challenge nonce to avoid sending user information and password between the client and the server.

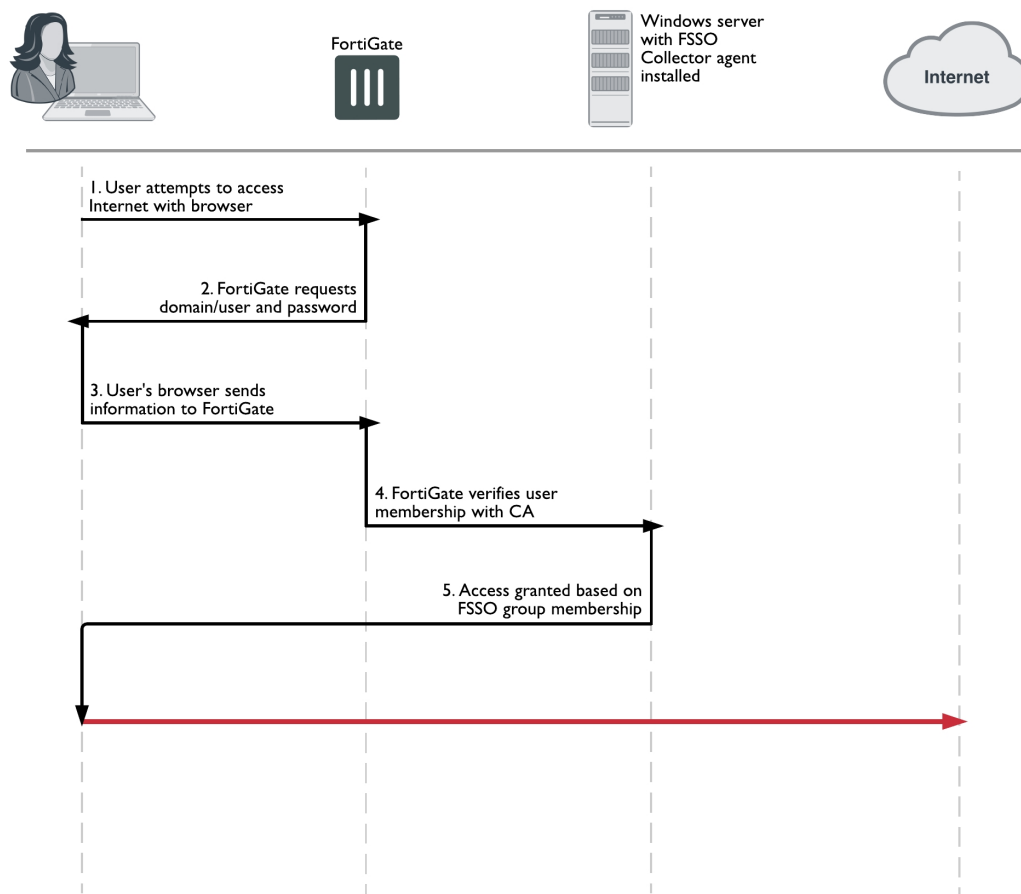
FSSO NTLM authentication support

In a Windows AD network, FSSO can also provide NTLM authentication service to the FortiGate unit. When the user makes a request that requires authentication, the FortiGate unit initiates NTLM negotiation with the client browser. The FortiGate unit does not process the NTLM packets itself. Instead, it forwards all the NTLM packets to the FSSO service to process.

NTLM has the benefit of not requiring an FSSO agent, but it is not transparent to users, and the user's web browser must support NTLM.

The NTLM protocol protects the user's password by not sending it over the network. Instead, the server sends the client a random number that the client must encrypt with the hash value of the user's password. The server compares the result of the client's encryption with the result of its own encryption. The two will match only if both parties used the same password.

NTLM authentication



If the NTLM authentication with the Windows AD network is successful, and the user belongs to one of the groups permitted in the applicable security policy, the FortiGate unit allows the connection but will require authentication again in the future when the current authentication expires.

Fortinet has tested NTLM authentication with Internet Explorer and Firefox browsers.

NTLM in a multiple domain environment

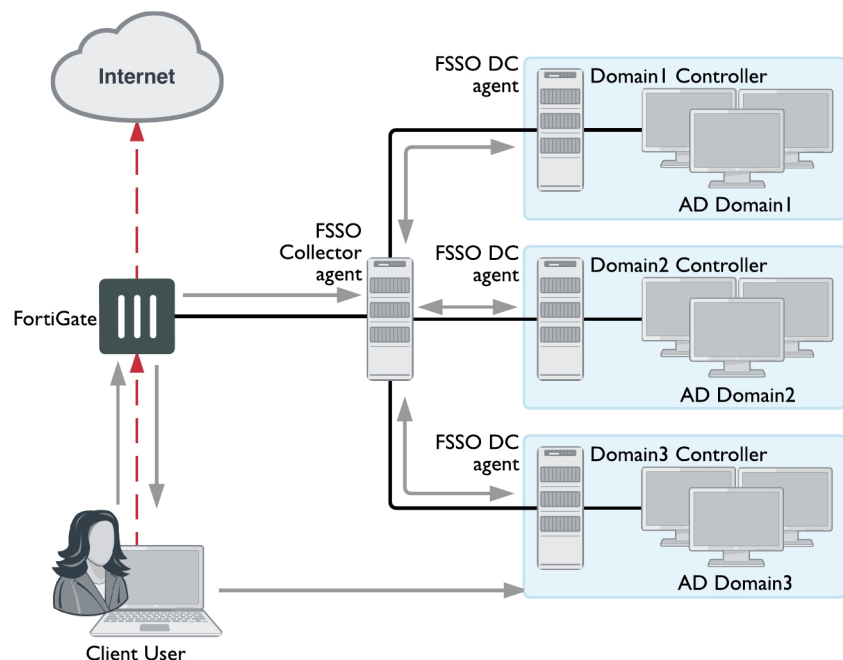
In a multiple domain environment for NTLM, the important factor is that there is a trust relation between the domains. In a forest, this relation is automatically created. So you can install FSSO agent on one of the domain controllers without worry.

But in case of multiple domains that are not in a forest, you need to create a trust relation between the domains. If you do not want to have a trust relation between your multiple domains, you need to use FSAE 4.0 MR1 and the DC agent needs to be installed once on each domain. Then you can use security policies to configure server access.

In the figure below, three domains are shown connected to the FSSO Collector agent server. The Client logs on to their local Domain Controller, which then sends the user logon event information to the Collector Agent. When the Client attempts to access the Internet, the FortiGate unit contacts the Collector Agent for the logon information, sees the Client is authenticated, and allows access to the Internet. There are multiple domains each

with a domain controller agent (DCagent) that sends logon information to the Collector agent. If the multiple domains have a trust relationship, only one DCagent is required instead of one per domain.

FSSO NTLM with multiple domains not in a forest



Understanding the NTLM authentication process

1. The user attempts to connect to an external (internet) HTTP resource. The client application (browser) on the user's computer issues an unauthenticated request through the FortiGate unit.
2. The FortiGate is aware that this client has not authenticated previously, so responds with a 401 Unauthenticated status code, and tells the client which authentication method to reply with in the header: Proxy-Authenticated: NTLM. Then the initial session is dismantled.
3. The client application connects again to the FortiGate, and issues a GET-request, with a Proxy-Authorization: NTLM <negotiate string> header. <negotiate-string> is a base64-encoded NTLM Type 1 negotiation packet.
4. The FortiGate unit replies with a 401 "proxy auth required" status code, and a Proxy-Authenticate: NTLM <challenge string> (a base 64-encoded NTLM Type 2 challenge packet). In this packet is the challenge nonce, a random number chosen for this negotiation that is used once and prevents replay attacks.



The TCP connection must be kept alive, as all subsequent authentication-related information is tied to the TCP connection. If it is dropped, the authentication process must start again from the beginning.

5. The client sends a new GET-request with a header: Proxy-Authenticate: NTLM <authenticate string>, where <authenticate string> is a NTLM Type 3 Authentication packet that contains:

- username and domain
 - the challenge nonce encoded with the client password (it may contain the challenge nonce twice using different algorithms).
6. If the negotiation is successful and the user belongs to one of the groups permitted in the security policy, the connection is allowed. Otherwise, the FortiGate unit denies the authentication by issuing a 401 return code and prompts for a username and password. Unless the TCP connection is broken, no further credentials are sent from the client to the proxy.



If the authentication policy reaches the authentication timeout period, a new NTLM handshake occurs.

Agentless NTLM support

Agentless NTLM authentication can be configured directly from the FortiGate to the Domain Controller via SMB protocol (no agent is required).

Note that this authentication method is only supported for proxy policies.

Syntax

Note that `domain-controller` is only available when `method` is set to `ntlm` and/or `negotiate-ntlm` is set to `enable`.

```
config authentication scheme
  edit <name>
    set method ntlm
    set domain-controller <dc-setting>
  next
end

config user domain-controller
  edit <name>
    set ip-address <dc-ip>
    set port <port> - default = 445
    set domain-name <dns-name>
    set ldap-server <name>
  next
end
```

Agent installation

After reading the appropriate sections of [Introduction to agent-based FSSO on page 276](#) to determine which FSSO agents you need, you can proceed to perform the necessary installations.

Ensure you have administrative rights on the servers where you are installing FSSO agents. It is best practice to install FSSO agents using the built-in local administrator account. Optionally, you can install FSSO without an admin account. See [Installing FSSO without using an administrator account on page 288](#).



In Windows 2008 by default, you do not have administrative user rights if you are logged on as a user other than as the built-in administrator, even if you were added to the local Administrators group on the computer.

The FSSO installer first installs the Collector agent. You can then continue with installation of the DC agent, or you can install it later by going to **Start > Programs > Fortinet > Fortinet Single Sign On Agent > Install DC Agent**. The installer will install a DC agent on the domain controllers of all of the trusted domains in your network.



Each domain controller connection needs a minimum guaranteed 64kpbs bandwidth to ensure proper FSSO functionality. Traffic shapers configured on the FortiGate can help guarantee these minimum bandwidths.

Collector agent installation

To install FSSO, you must obtain the FSSO_Setup file from the [Fortinet Support web site](#). This is available as either an executable (.exe) or a Microsoft Installer (.msi) file. Then you follow these two installation procedures on the server that will run the Collector agent. This can be any server or domain controller that is part of your network. These procedures also install the DC Agent on all of the domain controllers in your network.



IPv6 source addresses can be configured for connecting to an FSSO agent.

To install the Collector agent:

1. Create an account with administrator privileges and a password that does not expire. See Microsoft Advanced Server documentation for help with this task.
To use a non-admin read only account, see [Installing FSSO without using an administrator account on page 288](#).
2. Log on to the account that you created in Step 1.
3. Double-click the `FSSOSetup.exe` file.
The Fortinet SSO Collector Agent Setup Wizard starts.
4. Select **Next**.
5. Read and accept the license agreement. Select **Next**.
6. Optionally, you can change the installation location. Select **Next**.
7. Optionally, change the **User Name**.
8. By default, the agent is installed using the currently running account. If you want FSSO to use another existing admin account, change the **User Name** using the format `DomainName \ UserName`. For example if the account is `jsmith` and the domain is **example_corp** you would enter `example_corp\jsmith`.
9. In the **Password** field, enter the password for the account listed in the **User Name** field.
10. Select **Next**.
11. Enable as needed:
 - Monitor user logon events and send the information to the FortiGate unit
 - Serve NTLM authentication requests coming from FortiGateBy default, both methods are enabled. You can change these options after installation.
12. Select the access method to use for Windows Directory:
13. Select **Standard** to use Windows domain and username credentials.
14. Select **Advanced** if you will set up LDAP access to Windows Directory.
See [Collector agent AD access mode - standard versus advanced on page 280](#).

15. Select **Next** and then select **Install**.

If you want to use DC Agent mode, ensure that **Launch DC Agent Install Wizard** is selected. This will start DC agent installation immediately after you select **Finish**.

16. Select **Finish**.



If you see an error such as Service Fortinet Single Sign On agent (service_FSAE) failed to start, there are two possible reasons for this. Verify the user account you selected has sufficient privileges to run the FSSO service. Also verify the computer system you are attempting to install on is a supported operating system and version.

DC agent installation

The FSSO_Setup file contains both the Collector agent and DC Agent installers, but the DC Agent installer is also available separately as either a .exe or .msi file named DCAgent_Setup.

To install the DC Agent

1. If you have just installed the Collector agent, the FSSO - Install DC Agent wizard starts automatically. Otherwise, go to **Start > Programs > Fortinet > Fortinet Single Sign On Agent > Install DC Agent**.
2. Select **Next**.
3. Read and accept the license agreement. Select **Next**.
4. Optionally, you can change the installation location. Select **Next**.
5. Enter the **Collector agent IP address**.
6. If the Collector agent computer has multiple network interfaces, ensure that the one that is listed is on your network. The listed **Collector agent listening port** is the default. Only change this if the port is already used by another service.
7. Select **Next**.
8. Select the domains to monitor and select **Next**.
9. If any of your required domains are not listed, cancel the wizard and set up the proper trusted relationship with the domain controller. Then run the wizard again by going to **Start > Programs > Fortinet > Fortinet Single Sign On Agent > Install DC Agent**.
10. Optionally, select users that you do not want monitored. These users will not be able to authenticate to FortiGate units using FSSO. You can also do this later. See [Configuring the FSSO collector agent for Windows AD on page 290](#).
11. Select **Next**.
12. Optionally, clear the check boxes of domain controllers on which you do not want to install the DC Agent.
13. Select the **Working Mode** as DC Agent Mode. While you can select Polling Mode here, in that situation you would not be installing a DC Agent. For more information, see [DC Agent mode on page 279](#) and [Polling mode on page 279](#).
14. Select **Next**.
15. Select **Yes** when the wizard requests that you reboot the computer.



If you reinstall the FSSO software on this computer, your FSSO configuration is replaced with default settings.

If you want to create a redundant configuration, repeat the Collector agent installation procedure on at least one other Windows AD server.



When you start to install a second Collector agent, cancel the Install Wizard dialog appears the second time. From the configuration GUI, the monitored domain controller list will show your domain controllers un-selected. Select the ones you wish to monitor with this Collector agent, and select **Apply**.

Before you can use FSSO, you need to configure it on both Windows AD and on the FortiGate units. [Configuring FSSO on FortiGate units on page 310](#) will help you accomplish these two tasks.

Installing FSSO without using an administrator account

Normally when installing services in Windows, it is best to use the Domain Admin account, as stated earlier. This ensures installation goes smoothly and uninterrupted, and when using the FSSO agent there will be no permissions issues. However, it is possible to install FSSO with a non-admin account in Windows 2003 or 2008 AD.



The following instructions for Windows 2003 are specific to the event log polling mode only. Do not use this procedure with other FSSO configurations.

Windows 2003

There are two methods in Windows 2003 AD for installing FSSO without an admin account — add the non-admin user to the security log list, and use a non-admin account with read-only permissions. A problem with the first method is that full rights (read, write, and clear) are provided to the event log. This can be a problem when audits require limited or no write access to logs. In those situations, the non-admin account with read-only permissions is the solution.

To add the non-admin user account to the Windows 2003 security log list :

1. Go to **Default Domain Controller Security Settings > Security Settings > User Rights Assignment > Manage auditing and security log**.
2. Add the user account to this list.
3. Repeat these steps on every domain controller in Windows 2003 AD.
A reboot is required.

To use a non-admin account with read-only permissions to install FSSO on Windows 2003:

The following procedure provides the user account specified with read only access to the Windows 2003 AD Domain Controller Security Event Log which allows FSSO to function.

1. Find out the SID of the account you intend to use.
Tools for this can be downloaded for free from <http://technet.microsoft.com/en-us/sysinternals/bb897417>.
2. Then create the permission string. For example:
 - (A;;0x1;;;S-1-5-21-4136056096-764329382-1249792191-1107)
 - A means Allow,

- 0x1 means Read, and
 - S-1-5-21-4136056096-764329382-1249792191-1107 is the SID.
3. Then, append it to the registry key:
HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Eventlog\Security\CustomSD
 4. Repeat these steps on every domain controller in Windows 2003 AD.
A reboot is required.

Windows 2008

In Windows 2008 AD, if you do not want to use the Domain Admin account then the user account that starts the FSSO agent needs to be added to the Event Log Readers group.

When the user is added to the Event Log Readers group, that user is now allowed to have read only access to the event log and this is the minimal rights required for FSSO to work.

Citrix TS agent installation

To install the Citrix TS agent, you must obtain the TSAgent_Setup file from the [Fortinet Support web site](#). Perform the following installation procedure on the Citrix server.

To install the FSSO TS agent:

1. On the Citrix server, create an account with administrator privileges and a password that does not expire. See Citrix documentation for more information.
2. Log on to the account that you created in step 1.
3. Double-click the TSAgent_Setup installation file.
The Fortinet SSO Terminal Server Agent Setup Wizard starts.
4. Select **Next**.
5. Read and accept the license agreement. Select **Next**.
6. Optionally, you can change the installation location. Select **Next**.
7. Verify that **This Host IP Address** is correct.
8. In the **FSSO Collector Agent List**, enter the IP address(es) of your Collector Agents.
9. Select **Next** and then select **Install**.
The TS agent is installed.
10. Select **Finish**.

Novell eDirectory agent installation

To install the eDirectory agent, you must obtain the FSSO_Setup_eDirectory file from the [Fortinet Support web site](#). Perform the following installation procedure on the computer that will run the eDirectory agent. This can be any server or domain controller that is part of your network. You will need to provide some setup information.

To install the FSSO eDirectory agent:

1. Create an account with administrator privileges and a password that does not expire. See Novell documentation for more information.
2. Log on to the account that you created in step 1.
3. Double-click the FSSO_Setup_edirectory file to start the installation wizard.
4. Select **Next**.

5. Read and accept the license agreement. Select **Next**.
6. Optionally, change the installation location. Select **Next**.
7. Enter:

eDirectory Server	
Server Address	Enter the IP address of the eDirectory server.
Use secure connection (SSL)	Select to connect to the eDirectory server using SSL security.
Search Base DN	Enter the base Distinguished Name for the user search.

eDirectory Authentication	
Username	Enter a username that has access to the eDirectory, using LDAP format.
User password	Enter the password.

8. Select **Next**.
9. Select **Install**. When the installation completes, select **Finish**.

Updating FSSO agents on Windows AD

After FSSO is installed on your network, you may want to upgrade to a newer version. The following procedure helps ensure you have a trouble free upgrade. How you update FSSO depends on if you are using polling mode or DC Agent mode.

For polling mode, since there are no DC agents you only need to upgrade the Collector. However in DC Agent mode, each DC Agent must be updated as well.

To update FSSO in DC Agent mode:

1. Go to the system32 directory on all DC's and rename the `dcagent.dll` file to `dcagent.dll.old`. This ensures the when the upgrade is pushed to the DC it does not overwrite the old file. If there are any problems this makes it easy to revert to the old version.
2. Run the FSSO setup .exe file to update the collector. When this is completed, ignore any reboot message.
3. Go to **Programs > Fortinet > Fortinet Single Sign On Agent > Install DC Agent** and push the DC agent out to all servers. All DC's will now need to be rebooted so that the new DLL file is loaded.
4. After the reboot, go to all DC's and delete the `dcagent.dll.old` files.

Configuring the FSSO collector agent for Windows AD

On the FortiGate unit, security policies control access to network resources based on user groups. With Fortinet Single Sign On, this is also true but each FortiGate user group is associated with one or more Windows AD user groups. This is how Windows AD user groups get authenticated in the FortiGate security policy.

Fortinet Single Sign On sends information about Windows user logons to FortiGate units. If there are many users on your Windows AD domains, the large amount of information might affect the performance of the FortiGate units.

To avoid this problem, you can configure the Fortinet Single Sign On Collector agent to send logon information only for groups named in the FortiGate unit's security policies. See [Configuring FortiGate group filters on page 297](#).

On each server with a Collector agent, you will be

- [Configuring Windows AD server user groups](#)
- [Configuring collector agent settings](#), including the domain controllers to be monitored
- [Selecting Domain Controllers and working mode for monitoring](#)
- [Configuring directory access settings](#)
- [Configuring the ignore user list](#)
- [Configuring FortiGate group filters for each FortiGate unit](#)
- [Configuring FSSO ports](#)
- [Configuring alternate user IP address tracking](#)
- [Viewing FSSO component status](#)

Configuring Windows AD server user groups

FortiGate units control network resource access at the group level. All members of a user group have the same network access as defined in FortiGate security policies.

You can use existing Windows AD user groups for authentication to FortiGate units if you intend that all members within each group have the same network access privileges.

Otherwise, you need to create new user groups for this purpose.



If you change a user's group membership, the change does not take effect until the user logs off and then logs on again.



The FSSO Agent sends only Domain Local Security Group and Global Security Group information to FortiGate units. You cannot use Distribution group types for FortiGate access. No information is sent for empty groups.

Refer to Microsoft documentation for information about creating and managing Windows AD user groups.

Configuring collector agent settings

You need to configure which domain controllers the Collector agent will use and which domains to monitor for user logons. You can also alter default settings and settings you made during installation. These tasks are accomplished by configuring the FSSO Collector Agent, and selecting either Apply to enable the changes.

At any time to refresh the FSSO Agent settings, select Apply.

To configure the collector agent:

1. From the Start menu, select **Programs > Fortinet > Fortinet Single Sign-On Agent > Configure Fortinet Single Sign-On Agent**.
2. Enter the following information.

Monitoring user logon events	By default, this is enabled to automatically authenticate users as they log on to the Windows domain. Disable the Monitor feature only if you have a large network where this feature will slow responses too much.
Support NTLM authentication	By default, this is enabled to facilitate logon of users who are connected to a domain that does not have the FSSO DC Agent installed. Disable NTLM authentication only if your network does not support NTLM authentication for security or other reasons.
Collector Agent Status	Shows RUNNING when Collector agent is active.
Listening ports	You can change FSSO Collector Agent related port numbers if necessary.
FortiGate	TCP port for FortiGate units. Default 8000.
DC Agent	UDP port for DC Agents. Default 8002.
Logging	
Log level	Select the minimum severity level of logged messages.

Log file size limit (MB)	Enter the maximum size for the log file in MB. Default is 10.
View Log	View all Fortinet Single Sign On agent logs.
Log logon events in separate logs	<p>Record user login-related information separately from other logs. The information in this log includes:</p> <ul style="list-style-type: none"> • data received from DC agents • user logon/logoff information • workstation IP change information • data sent to FortiGate units
View Logon Events	If Log logon events in separate logs is enabled, you can view user login-related information.
Authentication	
Require authenticated connection from FortiGate	Select to require the FortiGate unit to authenticate before connecting to the Collector agent.
Password	Enter the password that FortiGate units must use to authenticate. The maximum password length is 15 characters. The default password is "fortinetcanada". It is highly recommended to modify this password.
Timers	
Workstation verify interval (minutes)	<p>Enter the interval in minutes at which the Fortinet Single Sign On Collector agent connects to client computers to determine whether the user is still logged on. The default is every 5 minutes. The interval may be increased if your network has too much traffic.</p> <p>Note: This verification process creates security log entries on the client computer.</p> <p>If ports 139 or 445 cannot be opened on your network, set the interval to 0 to prevent checking. See Configuring FSSO ports on page 299.</p>
Dead entry timeout interval	<p>Enter the interval in minutes after which Fortinet Single Sign On Agent purges information for user logons that it cannot verify. The default is 480 minutes (8 hours).</p> <p>Dead entries usually occur because the computer is unreachable (such as in standby mode or disconnected) but the user has not logged off. A common reason for this is when users forget to logoff before leaving the office for the day.</p> <p>You can also prevent dead entry checking by setting the interval to 0.</p>

IP address change verify interval	<p>Fortinet Single Sign On Agent periodically checks the IP addresses of logged-in users and updates the FortiGate unit when user IP addresses change. IP address verification prevents users from being locked out if they change IP addresses, as may happen with DHCP assigned addresses.</p> <p>Enter the verification interval in seconds. The default is 60 seconds. You can enter 0 to prevent IP address checking if you use static IP addresses.</p> <p>This does not apply to users authenticated through NTLM.</p>
Cache user group lookup result	<p>Enable caching.</p> <p>Caching can reduce group lookups and increase performance.</p>
Cache expire in (minutes)	<p>Fortinet Single Sign On Agent caches group information for logged-in users.</p> <p>Enter the duration in minutes after which the cache entry expires. If you enter 0, the cache never expires.</p> <p>A long cache expire interval may result in more stale user group information. This can be an issue when a user's group information is changed.</p>
Clear Group Cache	<p>Clear group information of logged-in users.</p> <p>This affects all logged-in users, and may force them to re-login.</p>

3. You can select **Save & close** now or leave the agent configuration window open to complete additional configuration in the following sections.



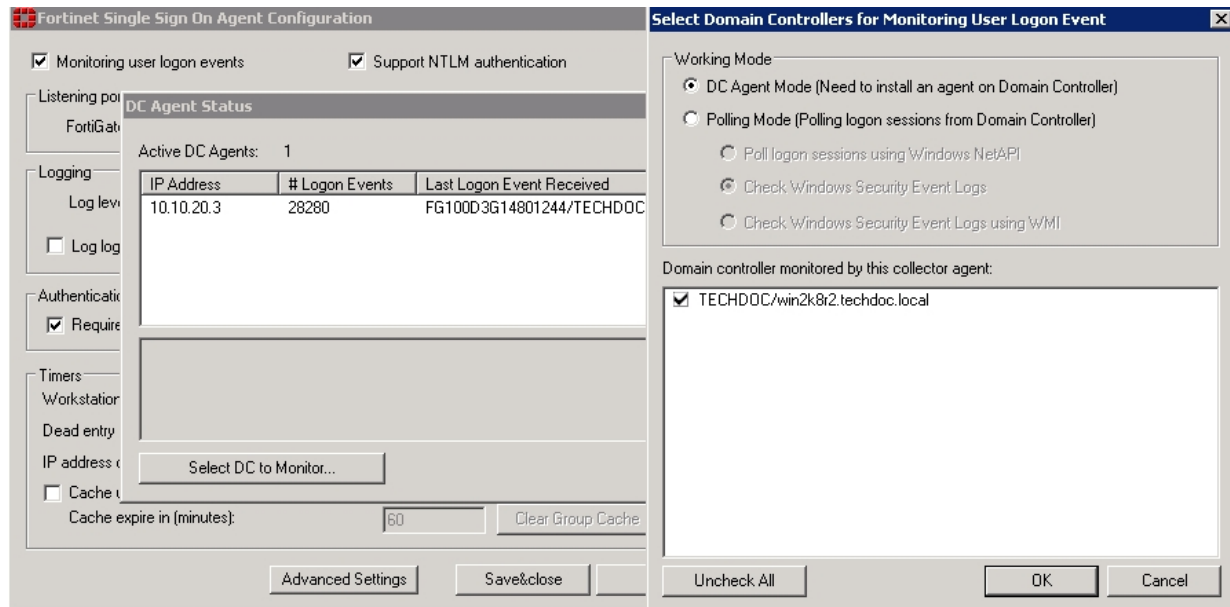
To view the version and build number information for your FSSO Collector Agent configuration, selecting the Fortinet icon in the upper left corner of the Collector agent Configuration screen and select **About Fortinet Single Sign On Agent configuration**.

Selecting Domain Controllers and working mode for monitoring

You can change which DC agents are monitored or change the working mode for logon event monitoring between DC agent mode and polling mode.

When polling mode is selected, it will poll port 445 of the domain controller every few seconds to see who is logged on.

1. From the Start menu select **Programs > Fortinet Fortinet Single Sign-On Agent > Configure Fortinet Single Sign On Agent**.
2. In the **Common Tasks** section, select **Show Monitored DCs**.
3. Select **Select DC to Monitor**.



4. Choose the **Working Mode**:

- **DC Agent mode** — a Domain Controller agent monitors user logon events and passes the information to the Collector agent. This provides reliable user logon information, however you must install a DC agent on every domain controller in the domain.
- **Polling mode** — the Collector agent polls each domain controller for user logon information. Under heavy system load this might provide information less reliably. However installing a DC agent on each domain controller is not required in this mode.

5. You also need to choose the method used to retrieve logon information:

- Poll logon sessions using Windows NetAPI
- Check Windows Security Event Logs
- Check Windows Security Event Logs using WMI

For more information about these options, see [Polling mode on page 279](#).

6. Select **OK**.

7. Select **Close**.

8. Select **Save & Close**.

Configuring directory access settings

The FSSO Collector Agent can access Windows Active Directory in one of two modes:

- **Standard** — the FSSO Collector Agent receives group information from the Collector agent in the **domain\user** format. This option is available on FortiOS 3.0 and later.
- **Advanced** — the FSSO Collector Agent obtains user group information using LDAP. The benefit of this method is that it is possible to nest groups within groups. This option is available on FortiOS 3.0 MR6 and later. The group information is in standard LDAP format.



If you change AD access mode, you must reconfigure your group filters to ensure that the group information is in the correct format.

To configure Directory Access settings:

1. From the Start menu select **Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent**.
2. In the **Common Tasks** section, select **Set Directory Access Information**.
The **Set Directory Access Information** dialog box opens.
3. From the **AD access mode** list, select either **Standard** or **Advanced**.
4. If you selected Advanced AD access mode, select **Advanced Setting** and configure the following settings and then select **OK**:

AD server address	Enter the address of your network's global catalog server.
AD server port	The default AD server port is 3268. This must match your server port.
BaseDN	Enter the Base distinguished name for the global catalog. This is the point in the tree that will be considered the starting point by default-See following example.
Username	If the global catalog accepts your Fortinet Single Sign On Agent agent's credentials, you can leave these fields blank. Otherwise, enter credentials for an account that can access the global catalog.
Password	

BaseDN example

An example DN for Training Fortinet Canada is `ou=training, ou=canada, dc=fortinet, dc=com`. If you set the **BaseDN** to `ou=canada, dc=fortinet, dc=com` then when Fortinet Single Sign On Agent is looking up user credentials, it will only search the Canada organizational unit, instead of all the possible countries in the company. Its a short cut to entering less information and faster searches.

However, you may have problems if you narrow the BaseDN too much when you have international employees from the company visiting different offices. If someone from Fortinet Japan is visiting the Canada office in the example above, their account credentials will not be matched because they are in `ou=japan, dc=fortinet, dc=com` instead of the BaseDN `ou=canada, dc=fortinet, dc=com`. The easy solution is to change the BaseDN to simply be `dc=fortinet, dc=com`. Then any search will check all the users in the company.

Configuring the ignore user list

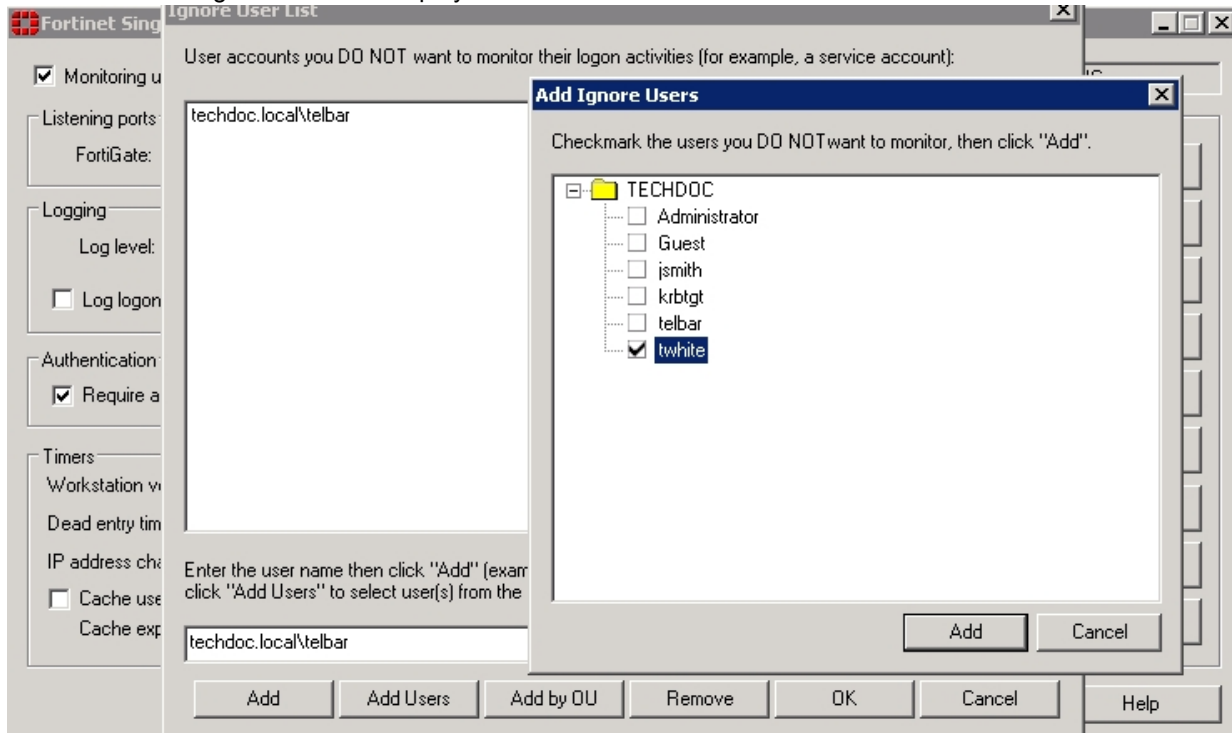
The ignore user list excludes users that do not authenticate to any FortiGate unit, such as system accounts. The logons of these users are not reported to FortiGate units. This reduces the amount of required resources on the FortiGate unit especially when logging logon events to memory.

To configure the ignore user list:

1. From the Start menu select **Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent**.

2. In the **Common Tasks** section, select **Set ignore user list**.

The current list of ignored users is displayed:



3. Do any of the following:

- To remove a user from the list, select the the username and then select **Remove**. The user's login is no longer ignored.
- To add users to be ignored,
 - enter the username in the format **domain\username** and select **Add** or
 - select **Add Users**, an **Add Ignore Users** window is displayed, checkmark the users you do not want to monitor, then select **Add** or
 - select **Add by OU**, an **Add Ignore Users by OU** window is displayed, select an OU from the directory tree, then select **Add**. All users under the selected OU will be added to the ignore user list.

4. Select **OK**.

Configuring FortiGate group filters

FortiGate group filters actively control which user logon information is sent to each FortiGate unit. You need to configure the group filter list so that each FortiGate unit receives the correct user logon information for the user groups that are named in its security policies. These group filters help limit the traffic sent to the FortiGate unit, and help limit the logon events logged.

The maximum number of Windows AD user groups allowed on a FortiGate depends on the model. Low end models support 256 Windows AD user groups, where mid and high end models support 1024 groups. This is per VDOM if VDOMs are enabled on the FortiGate unit.

You do not need to configure a group filter on the Collector agent if the FortiGate unit retrieves group information from Windows AD using LDAP. In that case, the Collector agent uses the list of groups you selected on the FortiGate unit as its group filter.

The filter list is initially empty. You need to configure filters for your FortiGate units using the Add function. At a minimum, create a default filter that applies to all FortiGate units without a defined filter.



If no filter is defined for a FortiGate unit and there is no default filter, the Collector agent sends all Windows AD group and user logon events to the FortiGate unit. While this normally is not a problem, limiting the amount of data sent to the FortiGate unit improves performance by reducing the amount of memory the unit uses to store the group list and resulting logs.

To configure a FortiGate group filter:

1. From the Start menu select **Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent**.
2. In the **Common Tasks** section, select **Set Group Filters**.
The FortiGate Filter List opens. It has the following columns:

FortiGate SN	The serial number of the FortiGate unit to which this filter applies.
Description	An optional description of the role of this FortiGate unit.
Monitored Groups	The Windows AD user groups that are relevant to the security policies on this FortiGate unit.
Add	Create a new filter.
Edit	Modify the filter selected in the list.
Remove	Remove the filter selected in the list.
OK	Save the filter list and exit.
Cancel	Cancel changes and exit.

3. Select **Add** to create a new filter. If you want to modify an existing filter, select it in the list and then select **Edit**.
4. Enter the following information and then select **OK**.

Default filter	Select to create the default filter. The default filter applies to any FortiGate unit that does not have a specific filter defined in the list.
FortiGate Serial Number	Enter the serial number of the FortiGate unit to which this filter applies. This field is not available if Default is selected.
Description	Enter a description of this FortiGate unit's role in your network. For example, you could list the resources accessed through this unit. This field is not available if Default is selected.
Monitor the following groups	The Collector agent sends to the FortiGate unit the user logon information for the Windows AD user groups in this list. Edit this list using the Add, Advanced and Remove buttons.

Add	<p>In the preceding single-line field, enter the Windows AD domain name and user group name, and then select Add. If you don't know the exact name, use the Advanced button instead.</p> <p>The format of the entry depends on the AD access mode (see Configuring directory access settings on page 295):</p> <p>Standard: Domain\Group</p> <p>Advanced: cn=group, ou=corp, dc=domain</p>
Advanced	Select Advanced , select the user groups from the list, and then select Add .
Remove	Remove the user groups selected in the monitor list.

Configuring FSSO ports

For FSSO to function properly a small number of TCP and UDP ports must be open through all firewalls on the network. These ports listed in this section assume the default FSSO ports are used.

TCP ports for FSSO agent with client computers

Windows AD records when users log on but not when they log off. For best performance, Fortinet Single Sign On Agent monitors when users log off. To do this, Fortinet Single Sign On Agent needs read-only access to each client computer's registry over TCP port 139 or 445. Open at least one of these ports — ensure it is not blocked by firewalls.

If it is not feasible or acceptable to open TCP port 139 or 445, you can turn off Fortinet Single Sign On Agent logoff detection. To do this, set the Collector agent **workstation verify interval** to 0. The FSSO Collector Agent assumes that the logged on computer remains logged on for the duration of the Collector agent dead entry timeout interval — by default this is eight hours.

Configuring ports on the collector agent computer

On the computer where you install the Collector agent, you must make sure that the firewall does not block the listening ports for the FortiGate unit and the DC Agent. By default, these are TCP port 8000 and UDP port 8002. For more information about setting these ports, see [Configuring FSSO advanced settings on page 305](#).

Configuring alternate user IP address tracking

In environments where user IP addresses change frequently, you can configure Fortinet Single Sign On Agent to use an alternate method to track user IP address changes. Using this method, Fortinet Single Sign On Agent responds more quickly to user IP address changes because it directly queries workstation IP addresses to match users and IP addresses.

This feature requires FSAE version 3.5.27 or later, Fortinet Single Sign On Agent any version, and FortiOS 3.0 MR7 or later.

To configure alternate user IP address tracking:

1. On the computer where the Collector agent is installed, go to **Start > Run**.
2. Enter `regedit` or `regedt32` and select **OK**.
The Registry Editor opens.
3. Find the registry key `HKEY_LOCAL_MACHINE\SOFTWARE\Fortinet\FSAE\collectoragent`.
4. Set the `supportFSAEauth` value (dword) to `00000001`.
If needed, create this new dword.
5. Close the Registry Editor.
6. From the Start menu select **Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent**.
7. Select **Apply**.
The Fortinet Single Sign On Agent service restarts with the updated registry settings.

Viewing FSSO component status

It is important to know the status of both your Collector agents and DC agents.

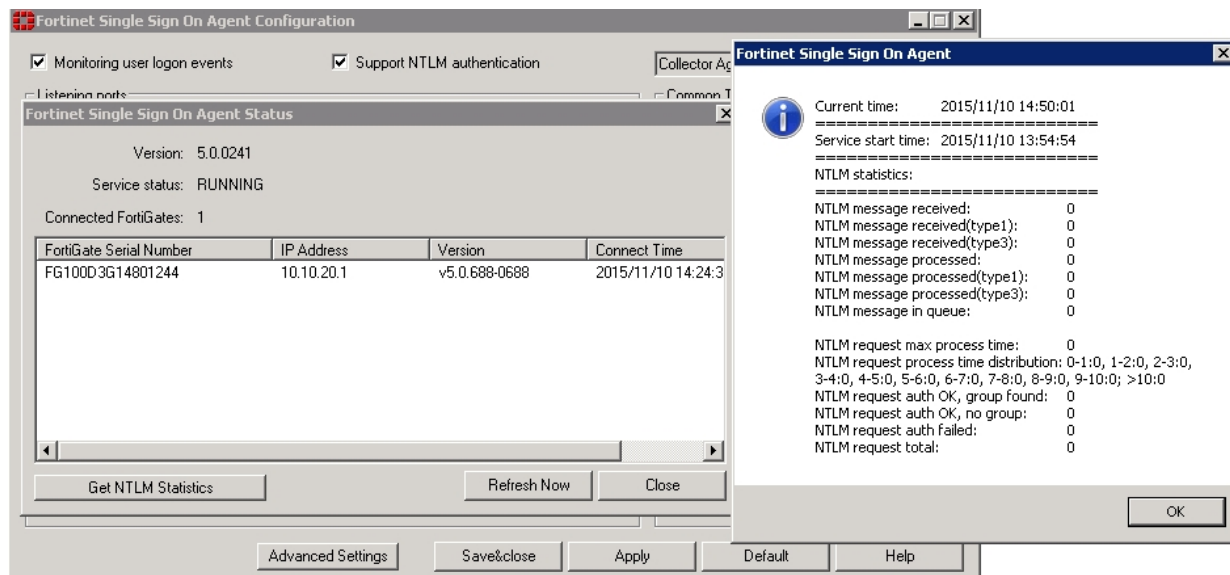
Viewing collector agent status

Use the **Show Service Status** to view your Collector agent information in the Status window. The Status window displays:

- the version of the software
- the status of the service
- the number of connected FortiGate units
- connected FortiGate information such as serial number, IP address, and connect time

To view Collector agent status:

1. From the Start menu select **Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent**.
2. In the **Common Tasks** section, select **Show Service Status**.
The Fortinet Single Sign On Collector agent Status window opens.
3. Optionally select **Get NTLM statistics** in the Status window to display NTLM information such as number of messages received, processed, failed, in the queue.



Viewing DC agent status

Use the **Show Monitored DCs** to view the status of DC agents.

To view domain controller agent status:

1. From the Start menu select **Programs > Fortinet > Fortinet Single Sign On Agent > Configure Fortinet Single Sign On Agent**.
2. In the **Common Tasks** section, select **Show Monitored DCs**.
For each DC Agent, the following information is displayed:
 - IP address
 - number of logon events received
 - the last logon event
 - when last logon was received

To change which DC agents are monitored or change the working mode for logon event monitoring, select **Select DC to Monitor**

DC Agent Status			
Active DC Agents: 1			
IP Address	# Logon Events	Last Logon Event Received	Received at
10.10.20.3	350	win2k8r2/KEEPALIVE/5.0.0241	2015/11/10 14:51:11
<div> <div>Select DC to Monitor...</div> <div>Refresh Now</div> <div>Close</div> </div>			

Configuring the FSSO TS agent for Citrix

The FSSO TS agent works with the same FSSO Collector agent that is used for integration with Windows Active Directory. Install the Collector agent first. Follow the Collector agent installation procedure in [Collector agent installation on page 286](#).



FortiOS 5.6 supports FSSO DC/TS Agent v.5.x.

Configuration steps include:

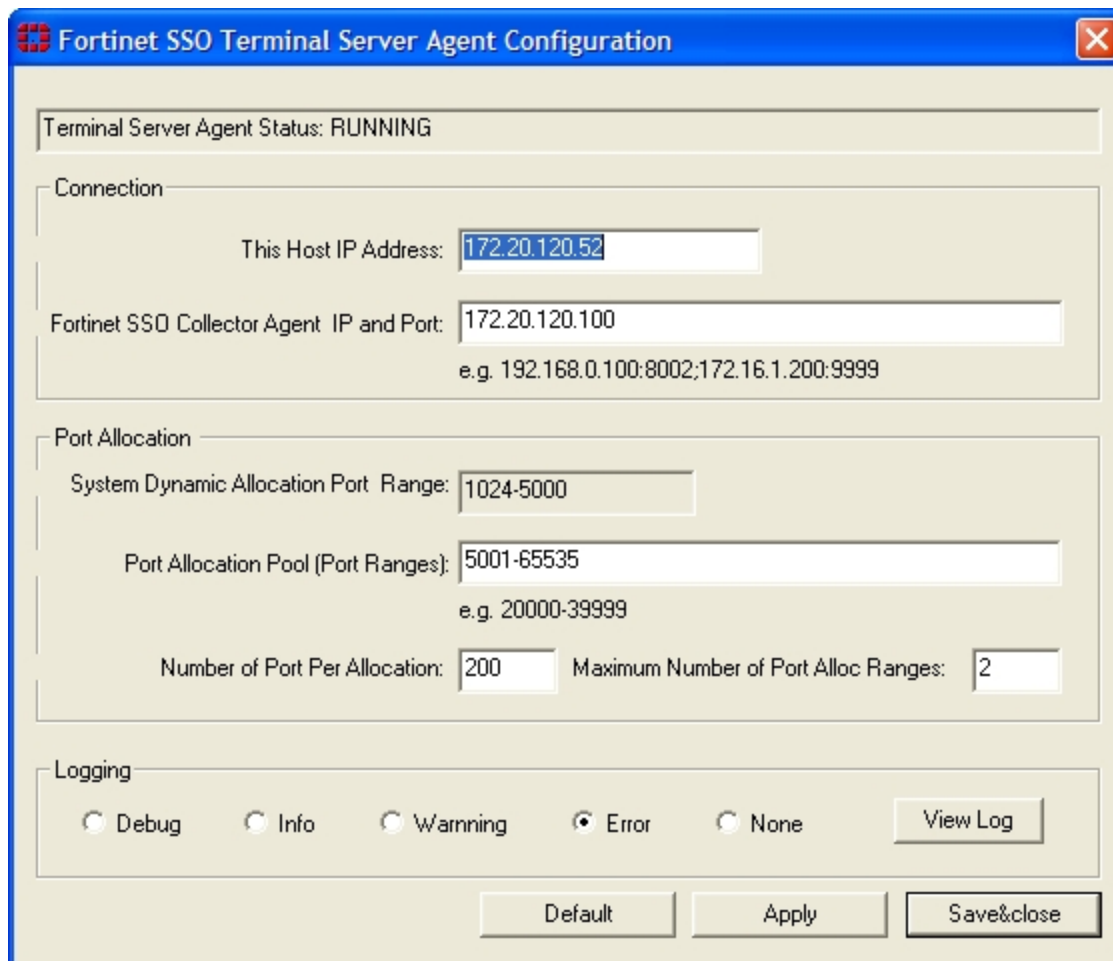
- Install the Fortinet Citrix FSSO agent on the Citrix server.
- Install the Fortinet FSSO collector on a server on the network.
- Add the Citrix FSSO agent to the FortiGate single-sign-On configuration.
- Add Citrix FSSO groups and users to an FSSO user group.
- Add an FSSO identity-based security policy that includes the Citrix FSSO user groups.

To change the TS agent configuration, select from the Start menu **Programs > Fortinet > Fortinet Single Sign-On Agent > TSAgent Config**. In addition to the host and Collector agent IP addresses that you set during installation, you can adjust port allocations for Citrix users. When a Citrix user logs on, the TS agent assigns that user a range of ports. By default each user has a range of 200 ports.



Fortinet SSO Collector Agent IP and Port needs to point to the current configured listening port on the collector which is port 8002 by default. Though it may be configured to a custom port.

Configuring the TS agent



The image shows the 'Fortinet SSO Terminal Server Agent Configuration' dialog box. It has a blue title bar with the Fortinet logo and a close button. The main area is divided into sections: 'Terminal Server Agent Status: RUNNING', 'Connection', 'Port Allocation', and 'Logging'. In the 'Connection' section, 'This Host IP Address' is set to '172.20.120.52' and 'Fortinet SSO Collector Agent IP and Port' is '172.20.120.100'. The 'Port Allocation' section shows 'System Dynamic Allocation Port Range' as '1024-5000', 'Port Allocation Pool (Port Ranges)' as '5001-65535', 'Number of Port Per Allocation' as '200', and 'Maximum Number of Port Alloc Ranges' as '2'. The 'Logging' section has radio buttons for 'Debug', 'Info', 'Warning', 'Error' (selected), and 'None', along with a 'View Log' button. At the bottom are 'Default', 'Apply', and 'Save&close' buttons.

Fortinet SSO Terminal Server Agent Configuration

Terminal Server Agent Status: RUNNING

Connection

This Host IP Address: 172.20.120.52

Fortinet SSO Collector Agent IP and Port: 172.20.120.100
e.g. 192.168.0.100:8002;172.16.1.200:9999

Port Allocation

System Dynamic Allocation Port Range: 1024-5000

Port Allocation Pool (Port Ranges): 5001-65535
e.g. 20000-39999

Number of Port Per Allocation: 200 Maximum Number of Port Alloc Ranges: 2

Logging

☐ Debug ☐ Info ☐ Warning ☒ Error ☐ None View Log

Default Apply Save&close

Configuring FSSO with Novell networks

You need to configure the eDirectory agent for it to communicate with eDirectory servers. You may have provided some of this information during installation.

This section includes:

- [Configuring the eDirectory agent](#)
- [Adding an eDirectory server](#)
- [Configuring a group filter](#)

Configuring the eDirectory agent

You need to configure the eDirectory agent for it to communicate with eDirectory servers.

To configure the eDirectory agent:

1. From the Start menu select **Programs > Fortinet > eDirectory Agent > eDirectory Config Utility**.
2. The eDirectory Agent Configuration Utility dialog opens. Enter the following information and select **OK**.

eDirectory Authentication

Username	Enter a username that has access to the eDirectory, using LDAP format.
Password	Enter the password.
Listening port	Enter the TCP port on which Fortinet Single Sign On Agent listens for connections from FortiGate units. The default is 8000. You can change the port if necessary.
Refresh interval	Enter the interval in seconds between polls of the eDirectory server to check for new logons. The default is 30 seconds.

FortiGate Connection Authentication

Require authenticated connection from FortiGate	Select to require the FortiGate unit to authenticate before connecting to the eDirectory Agent.
Password	Enter the password that FortiGate units must use to authenticate. The maximum password length is 15 characters. The default password is "FortinetCanada".
User logon Info Search Method	Select how the eDirectory agent accesses user logon information: LDAP or Native (Novell API). LDAP is the default. If you select Native , you must also have the Novell Client installed on the PC.

Logging

Log file size limit (MB)	Enter the maximum size for the log file in MB.
View Log	View the current log file.
Dump Session	List the currently logged-on users in the log file. This can be useful for troubleshooting.
Log level	Select Debug , Info , Warning or Error as the minimum severity level of message to log or select None to disable logging.

eDirectory Server List

Add	Add an eDirectory server. See Adding an eDirectory server on page 305 .
Delete	Delete the selected eDirectory server.
Edit	Modify the settings for the selected server.
Set Group Filters...	Select the user groups whose user logons will be reported to the FortiGate unit. This is used only if user groups are not selected on the FortiGate unit.

Adding an eDirectory server

Once the eDirectory agent is configured, you add one or more eDirectory servers.

To add an eDirectory server:

1. In the eDirectory Agent Configuration Utility dialog box (see the preceding procedure, [Configuring the eDirectory agent](#)), select **Add**.
2. The eDirectory Setup dialog box opens. Enter the following information and select OK:

eDirectory Server Address	Enter the IP address of the eDirectory server.
Port	If the eDirectory server does not use the default port 389, clear the Default check box and enter the port number.
Use default credential	Select to use the credentials specified in the eDirectory Configuration Utility. See Configuring the eDirectory agent on page 303 . Otherwise, leave the check box clear and enter a username and Password below.
User name	Enter a username that has access to the eDirectory, using LDAP format.
User password	Enter the password.
Use secure connection (SSL)	Select to connect to the eDirectory server using SSL security.
Search Base DN	Enter the base Distinguished Name for the user search.

Configuring a group filter

The eDirectory agent sends user logon information to the FortiGate unit for all user groups unless you either configure an LDAP server entry for the eDirectory on the FortiGate unit and select the groups that you want to monitor or configure the group filter on the eDirectory agent.

If both the FortiGate LDAP configuration and the eDirectory agent group filter are present, the FortiGate user group selections are used.

To configure the group filter:

1. From the Start menu select **Programs > Fortinet > eDirectory Agent > eDirectory Config Utility**.
2. Select **Set Group Filters**.
3. Do one of the following:
 - Enter group names, then select **Add**.
 - Select **Advanced**, select groups, and then select **Add**.
4. Select **OK**.

Configuring FSSO advanced settings

Depending on your network topologies and requirement, you may need to configure advanced settings in the FSSO Collector agent. To do so, from the Start menu, select **Programs > Fortinet > Fortinet Single Sign-On Agent > Configure Fortinet Single Sign-On Agent**, then from the **Common Tasks** section, select **Advanced Settings**.

This section include :

- [General settings](#)
- [Citrix/Terminal server](#)
- [Exchange server](#)
- [RADIUS accounting](#)

General settings

In the General tab, enter the following information and select **OK**.

Worker thread count	Number of threads started in the CA process. Default is 128 on CA version 5.0.0241.
Maximum FortiGate connections	Number of FortiGates can be connected to the CA. Default is 64.
Group look-up interval	The interval in seconds to lookup users/groups. If an AD group membership of currently logged on user, CA can detect this and update information on the FortiGate. Enter 0 for no checking.
Windows security Event logs	Choose the event logs to poll.
Event IDs to poll	<p>The default set (0) includes Kerberos authentication event logs : 672 for Windows server 2003, 4768 for Windows server 2008 and 2012 and NTLM authentication event logs : 680 for Windows server 2003, 4776 for Windows server 2008 and 2012. The extended set (1) includes Kerberos service ticket event logs : 673 for Windows server 2003, 4769 for Windows server 2008 and 2012. Service tickets are obtained whenever a user or computer accesses a server on the network.</p> <p>List the event ids separated by ",".</p>
Workstation Check	Optionally enable Use WMI to check user logoff for the collector agent to query whether users is still logged on.
Workstation Name Resolution Advance Options	
Alternative DNS server(s)	Collector Agent uses the DNS server configured on the machine it is running on by default. If CA should use another DNS server then one or more alternative DNS server can be configured here.
Alternative workstation suffix(es)	If only host name is available CA uses the default domain suffix to build a FQDN for DNS queries. In case CA should use a different suffix, it can be configured as well.

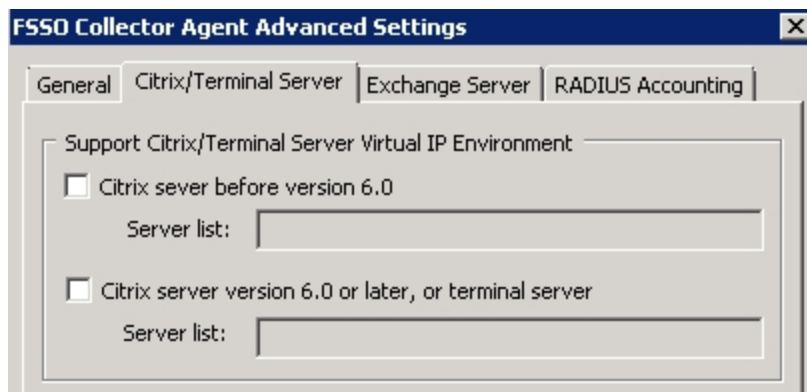
The screenshot shows the 'FSSO Collector Agent Advanced Settings' dialog box with the 'Citrix/Terminal Server' tab selected. The 'General' tab is also visible. The settings are as follows:

- Worker thread count: 128
- Maximum FortiGate connections: 64
- Group lookup interval (in seconds): 0 (0 for no checking)
- Windows Security Event Logs:
 - Event IDs to poll: 0 (0:default set, 1:extended set, or list the eventids separated by ';')
- Workstation Check:
 - ☒ Use WMI to check user logoff
- Workstation Name Resolution Advanced Options:
 - Alternative DNS server(s):
 - Alternative workstation suffix(es):

Citrix/Terminal server

In the Citrix/Terminal server tab, enter the following information and select **OK**.

Support Citrix/Terminal Server Virtual IP Environment	When Citrix server are configured with VIP, CA can get user logon events from these server. Citrix changed their interface and data format so version of Citrix server is important.
Citrix server before version 6.0	Enable this option if you Citrix server version is before 6.0.
Server list	Enter the list of servers separated by colon.
Citrix server version 6.0 or later, or Terminal Server	Enable this option if you Citrix server version is 6.0 or later.
Server list	Enter the list of servers separated by colon.

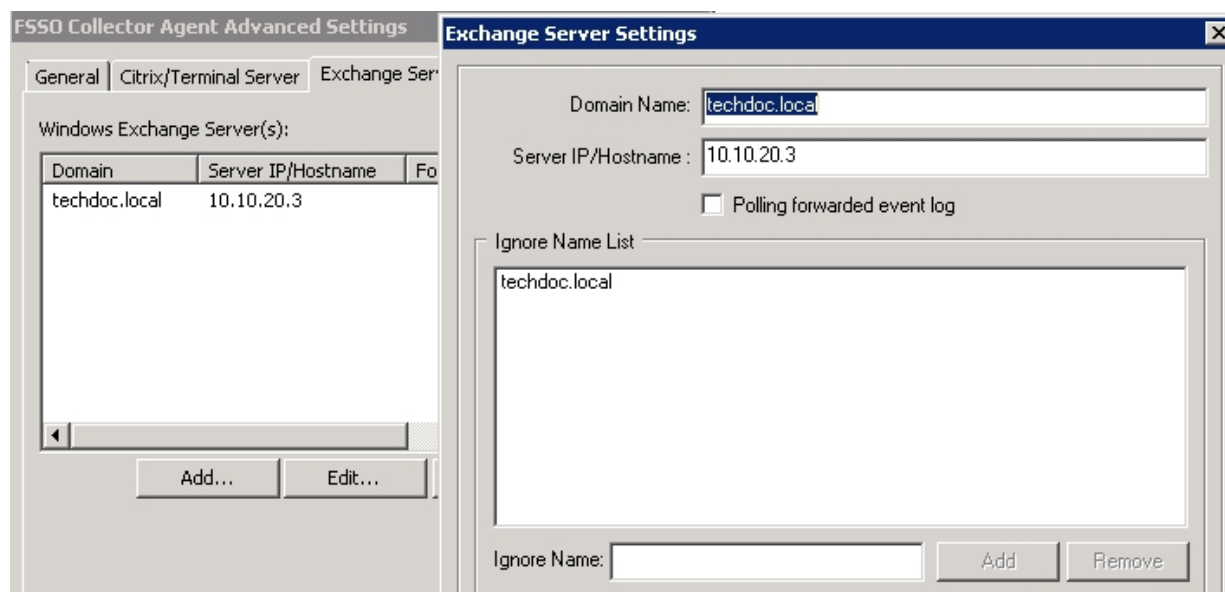


Exchange server

FSSO supports monitoring Microsoft Exchange server. This is useful for situation that the user use the domain account to access their email, but client device might or might not be in the domain. Support for Exchange server is configured on the Back-end FSSO collector agent under **Advanced Settings > Exchange Server**.

Select **Add** and enter the following information and select **OK**.

Domain Name	Enter your domain name.
Server IP/Hostname	Enter the IP address or the hostname of your exchange server.
Polling forwarded event log	This option for scenarios when you do not want that CA polls the Exchange Server logs directly. In this case you need to configure event log forwarding on the Exchange server. Exchange event logs can be forwarded to any member server. If you enable this, instead of the IP of the Exchange server configured in the previous step, you must then configure the IP of this member server. CA will then contact the member server.
Ignore Name	Because CA will also check Windows log files for logon events and when a user authenticates to Exchange Server there is also a logon event in Windows event log, which CA will read and this will overwrite the Exchange Server logon event (ES-EventLog) on CA. So it is recommended to set the ignore list to the domain the user belongs to. To do so, enter the domain name in the Ignore Name field and select Add .



RADIUS accounting

A RADIUS server must be configured in your network to send accounting messages to the Collector Agent which can be configured to work with most RADIUS-based accounting systems. In most cases, you only need to do the following to your RADIUS accounting system:

- Add a user group name field to customer accounts on the RADIUS server so that the name is added to the RADIUS Start record sent by the accounting system to the Collector Agent. User group names do not need to be added for all users, only to the accounts of users who will use RADIUS Accounting feature on the Collector Agent.
- Configure your accounting system to send RADIUS Start records to the Collector Agent.

The Collector Agent should be configured to listen for RADIUS accounting messages as following.


RADIUS Accounting Server	
Enable RADIUS Accounting Server	Enable this option to allow the CA to gather information about authenticated users via a RADIUS server and send these information to the FortiGate unit for monitoring.
Listen port	The port on which CA listens for RADIUS accounting messages. Default RADIUS accounting is 1813, but if RADIUS server sends accounting messages on different port, value can be configured here.
Shared secret	Common secret between CA and RADIUS server.

Default domain name

This should be the AD domain for which this CA is configured. In this case user name in RADIUS accounting message can be in simple format like user1.

If this value is empty, then user name in RADIUS accounting message must be in one of these formats user1@domain, Domain\user1 or domain/user1.

CA will use user name and domain to query group membership of user. Client IP address (Framed IP) should also be in RADIUS accounting message, so that CA can forward user name, IP address and groups to the FortiGate.



Configuring FSSO on FortiGate units

To configure your FortiGate unit to operate with agent-based FSSO, you

- Configure any access to LDAP servers that might be necessary. Skip this step if you are using FSSO Standard mode. See [Configuring LDAP server access on page 310](#).
- Specify the Collector agent or Novell eDirectory agent that will provide user logon information. See [Specifying your collector agents or Novell eDirectory agents on page 312](#).
- Add Active Directory user groups to FortiGate user groups. See [Creating FSSO user groups on page 313](#).
- Create security policies for FSSO-authenticated groups. See [Creating security policies on page 313](#).
- Optionally, specify a guest security policy to allow guest access. See [Enabling guest access through FSSO security policies on page 315](#).

Configuring LDAP server access

LDAP access is required if your network has a Novell eDirectory agent or a Collector agent using Windows Advanced AD access mode. If you are using FSSO Standard mode, go to [Specifying your collector agents or Novell eDirectory agents on page 312](#).

1. Go to **User & Device > LDAP Servers** and select **Create New**.
2. Enter the **Server IP/Name** and **Server Port** (default 389).
3. In the **Common Name Identifier** field, enter **sAMAccountName**. The default common name identifier is **cn**. This is correct for most LDAP servers. However some servers use other identifiers such as **uid**.
4. In the **Distinguished Name** field, enter your organization distinguished name. In this example, Distinguished

Name is dc=techdoc,dc=local

5. Select **Fetch DN**, this will fetch the Windows AD directory.

6. Set **Bind Type** to **Regular**.
7. In the **User DN** field, enter the administrative account name that you created for FSSO. For example, if the account is administrator, enter "administrator@techdoc.local".
8. Enter the administrative account password in the **Password** field.
9. Optionally select **Secure Connection**.
 - In the **Protocol** field, select **LDAPS** or **STARTTLS**.
 - In the **Certificate** field, select the appropriate certificate for authentication.

Note that you need to configure the Windows AD for secure connection accordingly.
10. Select **OK**.
11. Test your configuration by selecting the **Test** button. A successful message confirming the right settings appears.

To configure LDAP for FSSO - CLI example:

```
config user ldap
```

```

edit LDAP
    set server 10.10.20.3
    set cnid sAMAccountName
    set dn dc=techdoc,dc=local
    set type regular
    set username administrator@techdoc.local
    set password <your_password>
next
end

```

Specifying your collector agents or Novell eDirectory agents

You need to configure the FortiGate unit to access at least one Collector agent or Novell eDirectory agent. You can specify up to five servers on which you have installed a Collector or eDirectory agent. The FortiGate unit accesses these servers in the order that they appear in the list. If a server becomes unavailable, the next one in the list is tried.

To specify Collector agents - web-based manager:

1. Go to **Security Fabric > Fabric Connectors** and select **Create New**.
2. Under **SSO/Identity**, select **Fortinet Single Sign-On Agent**.
3. Enter a **Name** for the Windows AD server. This name appears in the list of Windows AD servers when you create user groups.
4. Enter the **Server IP/Name** and **Password** of the server where this agent is installed. Maximum name length is 63 characters. For the collector agent, passwords are only required only if you configured the agent to require authenticated access.
If the TCP port used for FSSO is not the default, 8000, you can change the setting in the CLI using the `config user fssso` command. See [Configuring collector agent settings on page 291](#).
5. Set **Collector Agent AD access mode** to **Advanced**, and select the **LDAP Server** you configured previously. See [Configuring LDAP server access on page 310](#).
6. Select the **Users** or **Groups** or **Organizational Units** tab to select the users, groups, OU that you want to monitor.
7. Select **OK**.

Name:
 Primary Agent IP/Name: Password:
 Secondary Agent IP/Name: Password: [More FSSO agents](#)
 LDAP Server: X
 Users/Groups

The screenshot shows the FortiGate web interface. On the left, the 'LDAP Tree' is expanded, showing the hierarchy: dc=techdoc,dc=local. On the right, the 'Users' tab is selected, displaying a table of users. The table has columns for ID, Name, and Full DN. The 'FortiOS Writers' user is highlighted with a green checkmark. At the bottom, a pagination bar shows '1 / 1' and 'Total: 38'.

ID	Name	Full DN
Domain Controllers	Domain Controllers	CN=Domain Controllers,CN=Users,DC=techdoc,DC=local
Domain Guests	Domain Guests	CN=Domain Guests,CN=Users,DC=techdoc,DC=local
Domain Users	Domain Users	CN=Domain Users,CN=Users,DC=techdoc,DC=local
Enterprise Admins	Enterprise Admins	CN=Enterprise Admins,CN=Users,DC=techdoc,DC=local
Enterprise Read-only Domain Controllers	Enterprise Read-only Domain Controllers	CN=Enterprise Read-only Domain Controllers,CN=Users,DC=techdoc,DC=local
Event Log Readers	Event Log Readers	CN=Event Log Readers,CN=Builtin,DC=techdoc,DC=local
FortiOS Writers	FortiOS Writers	CN=FortiOS Writers,CN=Users,DC=techdoc,DC=local
Group Policy Creator Owners	Group Policy Creator Owners	CN=Group Policy Creator Owners,CN=Users,DC=techdoc,DC=local
Guests	Guests	CN=Guests,CN=Builtin,DC=techdoc,DC=local
IIS_IUSRS	IIS_IUSRS	CN=IIS_IUSRS,CN=Builtin,DC=techdoc,DC=local

To specify the FSSO collector agent - CLI:

In this example, the SSO server name is techdoc and the LDAP server is LDAP.

```
config user fssso
  edit techdoc
    set ldap-server LDAP
    set password <your_password>
    set server 10.10.20.3
    set port 8000
  end
```

Creating FSSO user groups

You cannot use Windows or Novell groups directly in FortiGate security policies. You must create FortiGate user groups of the FSSO type and add Windows or Novell groups to them.

To create a user group for FSSO authentication - web-based manager:

1. Go to **User & Device > User Groups**.
2. Select **Create New**.
The New User Group dialog box opens.
3. In the **Name** box, enter a name for the group, FSSO_Internet_users for example.
4. In **Type**, select **Fortinet Single Sign-On (FSSO)**.
5. In **Members**, select the required **FSSO** groups.
6. Select **OK**.

To create the FSSO_Internet-users user group - CLI :

```
config user group
  edit FSSO_Internet_users
    set group-type fssso-service
    set member CN=Engineering,cn=users,dc=office,dc=example,dc=com
    CN=Sales,cn=users,dc=office,dc=example,dc=com
  end
```

Creating security policies

Policies that require FSSO authentication are very similar to other security policies. Using identity-based policies, you can configure access that depends on the FSSO user group. This allows each FSSO user group to have its own level of access to its own group of services

In this situation, Example.com is a company that has its employees and authentication servers on an internal network. The FortiGate unit intercepts all traffic leaving the internal network and requires FSSO authentication to access network resources on the Internet. The following procedure configures the security policy for FSSO authentication. FSSO is installed and configured including the RADIUS server, FSSO Collector agent, and user groups on the FortiGate

For the following procedure, the internal interface is `port1` and the external interface connected to the Internet is `port2`. There is an address group for the internal network called `company_network`. The FSSO user group is called `fssso_group`, and the FSSO RADIUS server is `fssso_rad_server`.

To configure an FSSO authentication security policy - web-based manager:

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information.

Incoming Interface	port1
Source Address	company_network
Source User(s)	fsso_group
Outgoing Interface	port2
Destination Address	all
Schedule	always
Service	HTTP, HTTPS, FTP, and Telnet
Action	ACCEPT
NAT	ON
UTM Security Profiles	ON for AntiVirus, IPS, Web Filter, and Email Filter, all using default profiles.
Log Allowed Traffic	ON. Select Security Events .

3. Select **OK**.
4. Ensure the FSSO authentication policy is higher in the policy list than more general policies for the same interfaces.

To create a security policy for FSSO authentication - CLI:

```
config firewall policy
  edit 0
    set srcintf port1
    set dstintf port2
    set srcaddr company_network
    set dstaddr all
    set action accept
    set groups fsso_group
    set schedule always
    set service HTTP HTTPS FTP TELNET
    set nat enable
  end
```

Here is an example of how this FSSO authentication policy is used. Example.com employee on the internal company network logs on to the internal network using their RADIUS username and password. When that user attempts to access the Internet, which requires FSSO authentication, the FortiGate authentication security policy intercepts the session, checks with the FSSO Collector agent to verify the user's identity and credentials, and then if everything is verified the user is allowed access to the Internet.

Users belonging to multiple groups

Before FSSO 4.0 MR3, if a user belonged to multiple user groups, the first security policy to match any group that user belonged to was the only security policy applied. If that specific group did not have access to this protocol or resource where another group did, the user was still denied access. For example, `test_user` belongs to `group1` and `group2`. There are two FSSO authentication policies — one matches `group1` to authenticate FTP traffic and one matches `group2` to authenticate email traffic. The `group1` policy is at the top of the list of policies. If `test_user` wants to access an email server, the first policy encountered for a group `test_user` belongs to is the `group1` policy which does not allow email access and `test_user` is denied access. This is despite the next policy allowing access to email. If the order was reversed in this case, the traffic would be matched and the user's traffic would be allowed through the firewall. However if the policy order was reversed, FTP traffic would not be matched.

As of FSSO 4.0 MR3, if a user belongs to multiple groups multiple then attempts to match the group are attempted if applicable. Using the above example, when the attempt to match the `group1` policy is made and fails, the next policy with a group that `test_user` is a member of is attempted. In this case, the next policy is matched and access is granted to the email server.

When configuring this example the only difference between the policies is the services that are listed and the FSSO user group name.

Authenticating through multiple groups allows administrators to assign groups for specific services, and users who are members of each group have access to those services. For example there could be an FTP group, an email group, and a Telnet group.

Enabling guest access through FSSO security policies

You can enable guest users to access FSSO security policies. Guests are users who are unknown to the Windows AD or Novell network and servers that do not logon to a Windows AD domain.

To enable guest access in your FSSO security policy, add an identity-based policy assigned to the built-in user group `SSO_Guest_Users`. Specify the services, schedule and protection profile that apply to guest users — typically guests receive reduced access to a reduced set of services.

IPv6 support for FSSO

FortiGate FSSO supports connecting to an FSSO agent over IPv6 and collecting and sending IPv6 details about endpoints. This is enforced in the same manner as IPv4 FSSO traffic.

Syntax

```
config user fsso
  edit <fsso agent name>
    set source-ip6 <IPv6 address for source>
  end
```

FortiOS FSSO log messages

There are two types of FortiOS log messages — firewall and event. FSSO-related log messages are generated from authentication events. These include user logon and log off events, and NTLM authentication events. These log messages are central to network accounting policies, and can also be useful in troubleshooting issues. For

more information on firewall logging, see "Enabling security logging". For more information on logging, see the FortiOS Handbook Log and Reporting guide.

Enabling authentication event logging

For the FortiGate unit to log events, that specific type of event must be enabled under logging.

When VDOMs are enabled certain options may not be available, such as CPU and memory usage events. You can enable event logs only when you are logged on to a VDOM; you cannot enable event logs globally.

To ensure you log all the events need, set the minimum log level to Notification or Information. Firewall logging requires Notification as a minimum. The closer to Debug level, the more information will be logged. While this extra information is useful, you must

To enable event logging:

1. Go to **Log & Report > Log Settings**.
2. In **Event Logging**, select:

System activity event	All system-related events, such as ping server failure and gateway status.
User activity event	All administration events, such as user logins, resets, and configuration updates.

3. Optionally you can enable any or all of the other logging event options.
4. Select **Apply**.

Authentication log messages

#	Date/Time	Level	Source	Action	Status	Message	Timestamp
1	06:15:24	Information	TELBAR (10.10.20.7)	FSSO-logout		FSSO-logout event from techdoc: user TELBAR logged off 10.10.20.7	11/12/2015, 6:15:24 AM
2	11-11 22:22	Information	ADMINISTRATOR (10.10.20.3)	authentication	logout	User ADMINISTRATOR succeeded in logout	11/11/2015, 10:22:15 P
3	11-11 22:22	Information	ADMINISTRATOR (10.10.20.3)	FSSO-logout		FSSO-logout event from techdoc: user ADMINISTRATOR logged off 10.10.20.3	11/11/2015, 10:22:15 P
4	11-11 22:17	Information	ADMINISTRATOR (10.10.20.3)	FSSO-logout		FSSO-logout event from techdoc: user ADMINISTRATOR logged on 10.10.20.3	11/11/2015, 10:17:12 P

1 / 5 [Total: 246]

#	1	Action	FSSO-logout
Date/Time	06:15:24	Dst	techdoc
Level	Information	Log Description	FSSO logout authentication status
Log ID	43015	Message	FSSO-logout event from techdoc: user TELBAR logged off 10.10.20.7
Source	TELBAR (10.10.20.7)	Sub Type	user
Timestamp	11/12/2015, 6:15:24 AM	User	TELBAR
Virtual Domain	root		

List of FSSO related log messages

Message ID	Severity	Description
43008	Notification	Authentication was successful
43009	Notification	Authentication session failed

Message ID	Severity	Description
43010	Warning	Authentication locked out
43011	Notification	Authentication timed out
43012	Notification	FSSO authentication was successful
43013	Notification	FSSO authentication failed
43014	Notification	FSSO user logged on
43015	Notification	FSSO user logged off
43016	Notification	NTLM authentication was successful
43017	Notification	NTLM authentication failed

For more information on logging, see the FortiOS Handbook Log and Reporting guide.

Using filters

Logon events are detected by the FSSO CA by monitoring the Security Event logs. Additional logon event filters, such as ServiceName and ServiceID, have been implemented so as to avoid instances of conflicting security events, where existing FSSO logon user information could be overwritten and impact user connectivity.

The problem arises when a scenario such as the following occurs:

1. **User1** logs on to **PC1** on 1.1.1.1, which is logged as a successful Kerberos logon event with an ID of 4769.
2. The FortiGate creates an authenticated FSSO user log entry for **User1/1.1.1.1**.
3. **User1** then maps a network drive and uses credentials for **User2** to logon to the same PC (**PC1**).
4. The FortiGate sees this as a separate logon to **PC1** by a new user, **User2**. As a result, the log entry is updated to **User2/1.1.1.1**.
5. If **User2** is a member of a different user group to **User1** (i.e. has different access permissions), **User1** could lose access to their network resources.

The new filter makes the CA ignore the event log created when **User1** mapped a network drive, meaning that the original entry for **User1** will not be changed.



While this is a useful feature, it is highly recommended to test filters in a redundant test environment.

Testing FSSO

Once FSSO is configured, you can easily test to ensure your configuration is working as expected. For additional FSSO testing, see [Troubleshooting FSSO on page 318](#).

1. Logon to one of the stations on the FSSO domain, and access an Internet resource.
2. Connect to the CLI of the FortiGate unit, and if possible log the output.
3. Enter the following command:

```
diagnose debug authd fsso list
```

4. Check the output. If FSSO is functioning properly you will see something similar to the following:

```
----FSSO logons----
IP: 10.10.20.3 User: ADMINISTRATOR Groups: CN=FORTIOS WRITERS,CN=USERS,DC=TECHDOC,DC=LOCAL
Workstation: WIN2K8R2.TECHDOC.LOCAL MemberOf: FortiOS_Writers
IP: 10.10.20.7 User: TELBAR Groups: CN=FORTIOS WRITERS,CN=USERS,DC=TECHDOC,DC=LOCAL
Workstation: TELBAR-PC7.TECHDOC.LOCAL
Total number of logons listed: 2, filtered: 0
----end of FSSO logons----
```

The exact information will vary based on your installation.

5. Check the FortiGate event log, for FSSO-auth action or other FSSO related events with FSSO information in the message field.
6. To check server connectivity, run the following commands from the CLI:

```
FGT# diagnose debug enable
FGT# diagnose debug authd fsso server-status
FGT# Server Name Connection Status Version
-----
techdoc          connected FSSO 5.0.0241
```

Troubleshooting FSSO

When installing, configuring, and working with FSSO some problems are quite common. A selection of these problems follows including explanations and solutions.

Some common Windows AD problems include:

- [General troubleshooting tips for FSSO](#)
- [Users on a particular computer \(IP address\) cannot access the network](#)
- [Guest users do not have access to network](#)
- [Can't find the DC agent service](#)
- [User logon events not received by FSSO collector agent](#)
- [Mac OS X users can't access external resources after waking from sleep mode](#)

General troubleshooting tips for FSSO

The following tips are useful in many FSSO troubleshooting situations.

- Ensure all firewalls are allowing the FSSO required ports through.
FSSO has a number of required ports that must be allowed through all firewalls or connections will fail. These include: ports 139, 389 (LDAP), 445, 636 (LDAP) 8000, and 8002.
- Ensure the Collector agent has at least 64kbps bandwidth to the FortiGate unit.
If not the Collector agent does not have this amount of bandwidth, information FSSO information may not reach the FortiGate unit resulting in outages. The best solution is to configure traffic shaping between the FortiGate unit and the Collector agent to ensure that minimum bandwidth is always available.

Users on a particular computer (IP address) cannot access the network

Windows AD Domain Controller agent gets the username and workstation where the logon attempt is coming from. If there are two computers with the same IP address and the same user trying to logon, it is possible for the

authentication system to become confused and believe that the user on computer_1 is actually trying to access computer_2.

Windows AD does not track when a user logs out. It is possible that a user logs out on one computer, and immediately logs onto a second computer while the system still believes the user is logged on the original computer. While this is allowed, information that is intended for the session on one computer may mistakenly end up going to the other computer instead. The result would look similar to a hijacked session.

Solutions

- Ensure each computer has separate IP addresses.
- Encourage users to logout on one machine before logging onto another machine.
- If multiple users have the same username, change the usernames to be unique.
- Shorten timeout timer to flush inactive sessions after a shorter time.

Guest users do not have access to network

A group of guest users was created, but they don't have access.

Solution

The group of the guest users was not included in a policy, so they do not fall under the guest account. To give them access, associate their group with a security policy.

Additionally, there is a default group called `SSO_Guest_Users`. Ensure that group is part of an identity-based security policy to allow traffic.

Can't find the DC agent service

The DCagent service can't be found in the list of regular windows services. This is because it has no associated Windows service.

Instead DCagent is really `dcagent.dll` and is located in the `Windows\system32` folder. This DLL file is loaded when windows boots up and it intercepts all logon events processed by the domain controller to send these events to the Collector agent (CA).

Solution

To verify that the DCagent is installed properly

1. Check that `DCagent.dll` exists in `Windows\system32` folder.
2. Check that the registry key exists: `[HKEY_LOCAL_MACHINE\SOFTWARE\Fortinet\FSAE\dcagent]`

If both exist, the DCagent is properly installed.

User logon events not received by FSSO collector agent

When a warning dialog is present on the screen on the Collector agent computer, the Collector agent will not receive any logon events. Once the dialog has been closed normal operation will resume.

If polling mode is enabled, it is possible the polling interval is too large. Use a shorter polling interval to ensure the collector agent is capturing all logon events.



The polling interval can only be adjusted using the firewall-embedded agent itself, not on agents deployed to member servers.

If NetAPI polling mode is enabled, consider switching to Event logs or Event Logs using WMI polling as it provides better accuracy.

Mac OS X users can't access external resources after waking from sleep mode

When client computers running Mac OS X (10.6.X and higher) wake up from sleep mode, the user must authenticate again to be able to access external resources. If the user does not re-authenticate, the user will maintain access to internal web sites, but will be unable to access any external resources.

This issue is caused by Mac OS X not providing sufficient information to the FSSO. This results in the FortiGate blocking access to the user because they cannot be authenticated.

Solution

The security settings on client computer(s) must be configured to require that a username and password be entered when exiting sleep mode or screen saver. With this feature enabled in Mac OS X, the FortiGate will receive the authentication information it requires to authenticate the user and allow them access.

Note that if the user reverts their settings to disable the password requirement, this will cause the issue to reappear.

SSO using RADIUS accounting records

A FortiGate unit can authenticate users transparently who have already authenticated on an external RADIUS server. Based on the user group to which the user belongs, the security policy applies the appropriate UTM profiles. RADIUS SSO is relatively simple because the FortiGate unit does not interact with the RADIUS server, it only monitors RADIUS accounting records that the server forwards (originating from the RADIUS client). These records include the user's IP address and user group.

After the initial set-up, changes to the user database, including changes to user group memberships, are made on the external RADIUS server, not on the FortiGate unit.

This section describes:

- [User's view of RADIUS SSO authentication](#)
- [Configuration overview](#)
- [Configuring the RADIUS server](#)
- [Creating the FortiGate RADIUS SSO agent](#)
- [Defining local user groups for RADIUS SSO](#)
- [Creating security policies](#)
- [Example - webfiltering for student and teacher accounts](#)

User's view of RADIUS SSO authentication

For the user, RADIUS SSO authentication is simple:

- The user connects to the RADIUS server and authenticates.
- The user attempts to connect to a network resource that is reached through a FortiGate unit. Authentication is required for access, but the user connects to the destination without being asked for logon credentials because the FortiGate unit knows that the user is already authenticated. FortiOS applies UTM features appropriate to the user groups that the user belongs to.

Configuration overview

The general steps to implement RADIUS Single Sign-On are:

1. If necessary, configure your RADIUS server. The user database needs to include user group information and the server needs to send accounting messages.
2. Create the FortiGate RADIUS SSO agent.
3. Define local user groups that map to RADIUS groups.
4. Create a security policy which specifies the user groups that are permitted access.

Configuring the RADIUS server

You can configure FortiGate RSSO to work with most RADIUS-based accounting systems. In most cases, you only need to do the following to your RADIUS accounting system:

- Add a user group name field to customer accounts on the RADIUS server so that the name is added to the RADIUS Start record sent by the accounting system to the FortiOS unit. User group names do not need to be added for all

users, only to the accounts of users who will use RSSO feature on the FortiGate unit.

- Configure your accounting system to send RADIUS Start records to the FortiOS unit. You can send the RADIUS Start records to any FortiGate network interface. If your FortiGate unit is operating with virtual domains (VDOMs) enabled, the RADIUS Start records must be sent to a network interface in the management VDOM.

IPv6 RADIUS support

RADIUS authentication is supported with IPv6, allowing administrators to configure an IPv6 RADIUS server on the FortiGate for IPv6 RADIUS authentication traffic to pass between the server and FortiGate.



Note that while you can set the primary RADIUS server's IPv6 address, the source IP address for communications to the RADIUS server cannot be configured as IPv6.

Syntax

Allow IPv6 access on an interface:

```
config system interface
  edit <name>
    config ipv6
      set ip6-allowaccess {ping | https | ssh | snmp | http | telnet | fgfm | capwap}
      set ip6-address <xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx/xxx>
    next
  next
end
```

Configure the IPv6 RADIUS server:

```
config user radius
  edit <name>
    set server <xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx:xxxx>
    ...
  next
end
```

Creating the FortiGate RADIUS SSO agent

Once you define a RADIUS SSO (RSSO) agent, the FortiGate unit will accept user logon information from any RADIUS server that has the same shared secret. You can create only one RSSO agent in each VDOM.

Before you create the RSSO agent, you need to allow RADIUS accounting information on the interface that connects to the RADIUS server.

To enable RADIUS access on the interface - web-based manager:

1. Go to **Network > Interfaces** and edit the interface to which the RADIUS server connected.
2. Select **Listen for RADIUS Accounting Messages**.
3. Select **OK**.

To enable RADIUS access on the interface - CLI:

In this example, the port2 interface is used.

```
config system interface
    edit port2
        set allowaccess radius-acct
    end
```

To create a RADIUS SSO agent:

1. Go to **[[[Undefined variable FortiOSGUIVariables.User & Device > Single Sign-On]]]** and select **Create New**.
2. In **Type**, select **RADIUS Single-Sign-On Agent**.
3. Select **Use RADIUS Shared Secret** and enter the RADIUS server shared secret.
4. Select **Send RADIUS Responses**.
5. Select **OK**.

To create a RADIUS SSO agent - CLI:

```
config user radius
    edit RSSO_Agent
        set rso enable
        set rso-validate-request-secret enable
        set rso-secret <your secret>
        set rso-radius-response enable
    end
```

Selecting which RADIUS attributes are used for RSSO

For RADIUS SSO to work, FortiOS needs to know the user's endpoint identifier (usually IP address) and RADIUS user group. There are default RADIUS attributes where FortiOS expects this information, but you can change these attributes in the `config user radius` CLI command.

RSSO information and RADIUS attribute defaults

RSSO Information	RADIUS Attribute	CLI field
Endpoint identifier	Calling-Station-ID	<code>rso-endpoint-attribute</code>
Endpoint block attribute	Called-Station-ID	<code>rso-endpoint-block-attribute</code>
User group	Class	<code>sso-attribute</code>
User	Prefix	<code>delegated-IPv6-prefix</code>
User	Prefix	<code>framed-IPv6-prefix</code>

The Endpoint block attribute can be used to block or allow a user. If the attribute value is set to the name of an attribute that indicates whether to block or allow, FortiOS blocks or allows respectively all traffic from that user's IP address. The RSSO fields are visible only when `rso` is set to `enable`.

The Prefix attributes allow for RSSO to provide a /56 prefix for DSL customers. All devices connected from the same location (/56 per subscriber) can be mapped to the same profile without the need to create multiple /64 or smaller entries.

Override SSO attribute

Prior to FortiOS 5.4, when receiving a new start message with a different group name for the same user, and a different IP address such as for a roaming mobile device, the original process was to override all group name information to the latest group name received from the latest start message.

You can disable this override when needed. The default behavior keeps the original design.

To enable or disable overriding SSO attribute - CLI

```
config user radius
edit <name>
set rso <enable>
set sso-attribute-value-override {enable | disable} Enable/disable override of old attribute value with new value for the same endpoint.
```

Configuring logging for RSSO

In the `config user radius` CLI command, you can set the following flags in the `rso-log-flags` field to determine which types of RSSO-related events are logged:

- `protocol-error` — A RADIUS protocol error occurred.
- `profile-missing` — FortiOS cannot find a user group name in a RADIUS start message that matches the name of an RSSO user group in FortiOS.
- `accounting-stop-missed` — a user context entry expired without FortiOS receiving a RADIUS Stop message.
- `accounting-event` — FortiOS did not find the expected information in a RADIUS record.
- `endpoint-block` — FortiOS blocked a user because the RADIUS record's endpoint block attribute had the value "Block".
- `radiusd-other` — Other events, described in the log message.

Defining local user groups for RADIUS SSO

You cannot use RADIUS user groups directly in security policies. Instead, you create locally-defined user groups on the FortiGate unit and associate each of them with a RADIUS user group.

To define local user groups for RADIUS SSO:

1. Go to **User & Device > User Groups** and select **Create New**.
2. Enter a Name for the user group.
3. In **Type**, select **RADIUS Single Sign-On (RSSO)**.
4. In **RADIUS Attribute Value**, enter the name of the RADIUS user group this local user group represents.
5. Select **OK**.

To define local user groups for RADIUS SSO:

This example creates an RSSO user group called RSSO-1 that is associated with RADIUS user group "student".

```
config user group
edit RSSO-1
set group-type rso
set sso-attribute-value student
end
```

Creating security policies

RADIUS SSO uses regular identity-based security policies. The RSSO user group you specify determines which users are permitted to use the policy. You can create multiple policies if user groups can have different UTM features enabled, different permitted services, schedules, and so on.

To create a security policy for RSSO - web-based manager:

1. Go to **Policy & Objects > IPv4 Policy**.
2. Select **Create New**.
3. Enter the following information.

Incoming Interface	as needed
Source Address	as needed
Source User(s)	Select the user groups you created for RSSO. See Defining local user groups for RADIUS SSO on page 324 .
Outgoing Interface	as needed
Destination Address	all
Schedule	as needed
Service	as needed
Action	ACCEPT
Enable NAT	Selected
Security Profiles	Select security profiles appropriate for the user group.

4. Select **OK**.

To ensure an RSSO-related policy is matched first, the policy should be placed higher in the security policy list than more general policies for the same interfaces.

5. Select **OK**.

To create a security policy for RSSO - CLI:

In this example, an internal network to Internet policy enables web access for members of a student group and activates the appropriate UTM profiles.

```
config firewall policy
edit 0
    set srcintf internal
    set dstintf wan1
    set srcaddr all
    set dstaddr "all"
    set action accept
    set rso enable
    set groups "RSSO-student"
```

```

set schedule always
set service HTTP HTTPS
set nat enable
set utm-status enable
set av-profile students
set webfilter-profile students
set spamfilter-profile students
set dlp-sensor default
set ips-sensor default
set application-list students
set profile-protocol-options "default"
end

```

Example - webfiltering for student and teacher accounts

The following example uses RADIUS SSO to apply web filtering to students, but not to teachers. Assume that the RADIUS server is already configured to send RADIUS Start and Stop records to the FortiGate unit. There are two RADIUS user groups, **students** and **teachers**, recorded in the default attribute **Class**. The workstations are connected to port1, port2 connects to the RADIUS server, and port3 connects to the Internet.

Configure the student web filter profile:

1. Go to **Security Profiles > Web Filter** and select **Create New** (the "+" button).
2. Enter the following and select **OK**.

Name	student
Inspection Mode	Proxy
FortiGuard Categories	Enable. Right-click the Potentially Liable category and select Block . Repeat for Adult/Mature Content and Security Risk .

Create the RADIUS SSO agent:

1. Go to **Security Fabric > Fabric Connectors** and select **Create New**.
2. Under **SSO/Identity**, select **RADIUS Single Sign-On Agent**.
3. Enter a name for the RSSO Agent.
4. Enable **Use RADIUS Shared Secret** and enter the RADIUS server's shared secret.
5. Enable **Send RADIUS Responses**.
6. Select **OK**.

Define local user groups associated with the RADIUS SSO user groups:

1. Go to **User & Device > User Groups** and select **Create New**.
2. Enter the following and select **OK**.

Name	RSSO-students
Type	RADIUS Single Sign-On (RSSO)
RADIUS Attribute Value	students

3. Select **Create New**, enter the following and select **OK**.

Name	RSSO-teachers
Type	RADIUS Single Sign-On (RSSO)
RADIUS Attribute Value	teachers

Create a security policy for students:

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter

Incoming Interface	port1
Source Address	all
Source User(s)	RSSO-students
Source Device Type	All
Outgoing Interface	port3
Destination Address	all
Schedule	always
Service	HTTP, HTTPS
Action	ACCEPT
NAT	ON
Security Profiles	Enable AntiVirus, Web Filter, IPS. In Web Filter, select the student profile.

3. Select **OK**.

Create a security policy for teachers:

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter

Incoming Interface	port2
Source Address	all
Source User(s)	RSSO-teachers
Source Device Type	All
Outgoing Interface	port3

Destination Address	all
Schedule	always
Service	ALL
Action	ACCEPT
NAT	ON
Security Profiles	Enable AntiVirus and IPS.

3. Select **OK**.

Monitoring authenticated users












This section describes how to view lists of currently logged-in firewall and VPN users. It also describes how to disconnect users.

The following topics are included in this section:

- [Monitoring firewall users](#)
- [Monitoring SSL VPN users](#)
- [Monitoring IPsec VPN users](#)
- [Monitoring users quarantine](#)

Monitoring firewall users

To monitor firewall users, go to **Monitor > Firewall User Monitor**.

 Refresh	 De-authenticate						
 User Name	 User Group	 Policy ID	 Duration	 IP Address	 Traffic Volume	 Method	
user3	Group1	2	0 day(s) 0 hour(s) 4 minute(s)	10.11.101.20	35 KB	FW-auth	
user4	Group1	2	0 day(s) 3 hour(s) 4 minute(s)	10.11.101.101	421 KB	FW-auth	

You can de-authenticate a user by selecting the Delete icon for that entry.


You can filter the list of displayed users by selecting the funnel icon for one of the column titles or selecting **Filter Settings**.

Optionally, you can de-authenticate multiple users by selecting them and then selecting **De-authenticate**.

Monitoring SSL VPN users

You can monitor web-mode and tunnel-mode SSL VPN users by username and IP address.

To monitor SSL VPN users, go to **Monitor > SSL-VPN Monitor**. To disconnect a user, select the user and then select the **Delete** icon.

<div> Delete</div>					
<input type="checkbox"/>	No.	User	Source IP	Begin Time	Description
<input type="checkbox"/>	1	user2	172.20.120.51	Wed Mar 17 13:17:32 2010	
<input checked="" type="checkbox"/>		Subsession			Tunnel IP:10.0.0.1

The first line, listing the username and IP address, is present for a user with either a web-mode or tunnel-mode connection. The Subsession line is present only if the user has a tunnel mode connection. The **Description** column displays the virtual IP address assigned to the user's tunnel-mode connection.

For more information about SSL VPN, see the FortiOS Handbook SSL VPN guide.

To monitor SSL VPN users - CLI:

To list all of the SSL VPN sessions and their index numbers:

```
execute vpn sslvpn list
```


The output looks like this:

```
SSL-VPN Login Users:
  Index   User   Auth Type   Timeout   From           HTTPS in/out
  0       user1  1           256      172.20.120.51  0/0

SSL-VPN sessions:
  Index   User   Source IP      Tunnel/Dest IP
  0       user2  172.20.120.51  10.0.0.1
```

You can use the Index value in the following commands to disconnect user sessions:

To disconnect a tunnel-mode user

```
execute vpn sslvpn del-tunnel <index>
```

To disconnect a web-mode user

```
execute vpn sslvpn del-web <index>
```

You can also disconnect multiple users:

To disconnect all tunnel-mode SSL VPN users in this VDOM

```
execute vpn ssl del-all tunnel
```

To disconnect all SSL VPN users in this VDOM

```
execute vpn ssl del-all
```

Monitoring IPsec VPN users

To monitor IPsec VPN tunnels in the web-based manager, go to **Monitor > IPsec Monitor**. user names are available only for users who authenticate with XAuth.

You can close a tunnel by selecting the tunnel and right click to select **Bring Down**.

Type	Dialup					
Name	Remote Gateway	Timeout	Status	Incoming Data	Outgoing Data	Username
dialup1_0	172.20.120.51	1116	 Bring Down	79233170 B	171639314 B	user2

For more information, see the FortiOS Handbook IPsec VPN guide.

Monitoring users quarantine

The user quarantine list shows all IP addresses and interfaces blocked by NAC quarantine. The list also shows all IP addresses, authenticated users, senders, and interfaces blocked by data leak prevention (DLP). The system administrator can selectively release users or interfaces from quarantine or configure quarantine to expire after a selected time period.

All sessions started by users or IP addresses on the user quarantine list are blocked until the user or IP address is removed from the list. All sessions to an interface on the list are blocked until the interface is removed from the list.

You can configure NAC quarantine to add users or IP addresses to the user quarantine list under the following conditions:

- **Users or IP addresses that originate attacks detected by IPS** - To quarantine users or IP addresses that originate attacks, enable and configure **Quarantine** in an IPS Filter.
- **Users or IP addresses that are quarantined by Data Leak Prevention** - In a DLP sensor select **Quarantine IP Address** as the action to take.

For more information, see [FortiOS Security Profiles](#) guide.

Users are viewed from **Monitor > Quarantine Monitor**.

Delete	Removes the selected user or IP address from the User Quarantine list.
Remove All	Removes all users and IP addresses from the User Quarantine list.
Search	Search the list for a particular IP address.
Source	The FortiGate function that caused the user or IP address to be added to the User Quarantine list: IPS or Data Leak Prevention.
Created	The date and time the user or IP address was added to the Banned User list.
Expires	The date and time the user or IP address will be automatically removed from the User Quarantine list. If Expires is Indefinite , you must manually remove the user or host from the list.

Examples and troubleshooting


This chapter provides an example of a FortiGate unit providing authenticated access to the Internet for both Windows network users and local users.

The following topics are included in this section:

- [Firewall authentication example](#)
- [LDAP dial-in using member-attribute example](#)
- [RADIUS SSO example](#)
- [Troubleshooting](#)

Firewall authentication example

Example configuration



Overview

In this example, there is a Windows network connected to Port 2 on the FortiGate unit and another LAN, Network_1, connected to Port 3.

All Windows network users authenticate when they logon to their network. Members of the Engineering and Sales groups can access the Internet without entering their authentication credentials again. The example assumes that the Fortinet Single Sign On (FSSO) has already been installed and configured on the domain controller.

LAN users who belong to the Internet_users group can access the Internet after entering their username and password to authenticate. This example shows only two users, User1 is authenticated by a password stored on the FortiGate unit, User2 is authenticated on an external authentication server. Both of these users are referred to as local users because the user account is created on the FortiGate unit.

Creating a locally-authenticated user account

User1 is authenticated by a password stored on the FortiGate unit. It is very simple to create this type of account.

To create a local user - web-based manager:

1. Go to **User & Device > User Definition** and select **Create New**.
2. Follow the User Creation Wizard, entering the following information and then select **Create**:

User Type	Local User
User Name	User1
Password	hardtoguess
Email Address SMS	(optional)
Enable	Select.

To create a local user - CLI:

```
config user local
edit user1
set type password
set passwd hardtoguess
end
```

Creating a RADIUS-authenticated user account

To authenticate users using an external authentication server, you must first configure the FortiGate unit to access the server.

To configure the remote authentication server - web-based manager:

1. Go to **User & Device > RADIUS Servers** and select **Create New**.
2. Enter the following information and select **OK**:

Name	OurRADIUSsrv
Primary Server Name/IP	10.11.101.15
Primary Server Secret	OurSecret
Authentication Scheme	Select Use Default Authentication Scheme .

To configure the remote authentication server - CLI:

```
config user radius
edit OurRADIUSsrv
set server 10.11.102.15
set secret OurSecret
set auth-type auto
end
```

Creation of the user account is similar to the locally-authenticated account, except that you specify the RADIUS authentication server instead of the user's password.

To configure a remote user - web-based manager:

1. Go to **User & Device > User Definition** and select **Create New**.
2. Follow the User Creation Wizard, entering the following information and then select **Create**:

User Type	Remote RADIUS User
User Name	User2
RADIUS server	OurRADIUSsrv
Email Address SMS	(optional)
Enable	Select

To configure a remote user - CLI:

```
config user local
  edit User2
    set name User2
    set type radius
    set radius-server OurRADIUSsrv
  end
```

Creating user groups

There are two user groups: an FSSO user group for FSSO users and a firewall user group for other users. It is not possible to combine these two types of users in the same user group.

Creating the FSSO user group

For this example, assume that FSSO has already been set up on the Windows network and that it uses Advanced mode, meaning that it uses LDAP to access user group information. You need to

- configure LDAP access to the Windows AD global catalog
- specify the collector agent that sends user logon information to the FortiGate unit
- select Windows user groups to monitor
- select and add the Engineering and Sales groups to an FSSO user group

To configure LDAP for FSSO - web-based manager:

1. Go to **User & Device > LDAP Servers** and select **Create New**.
2. Enter the following information:

Name	ADserver
Server Name / IP	10.11.101.160
Distinguished Name	dc=office,dc=example,dc=com

Bind Type	Regular
User DN	cn=FSSO_Admin,cn=users,dc=office,dc=example,dc=com
Password	Enter a secure password.

3. Leave other fields at their default values.
4. Select **OK**.

To configure LDAP for FSSO - CLI"

```
config user ldap
  edit "ADserver"
    set server "10.11.101.160"
    set dn "cn=users,dc=office,dc=example,dc=com"
    set type regular
    set username "cn=administrator,cn=users,dc=office,dc=example,dc=com"
    set password set_a_secure_password
  next
end
```

To specify the collector agent for FSSO - web-based manager:

1. Go to **Security Fabric > Fabric Connectors** and select **Create New**.
2. Under **SSO/Identity**, select **Fortinet Single Sign-On Agent**.
3. Enter a **Name** (in this example, WinGroups) for the Windows AD server. This name appears in the list of Windows AD servers when you create user groups.
4. Enter the **Server IP/Name** (in this example, 10.11.101.160) and **Password** (in this example, fortinet_canada) of the server where this agent is installed. Maximum name length is 63 characters. For the collector agent, passwords are only required only if you configured the agent to require authenticated access.
If the TCP port used for FSSO is not the default, 8000, you can change the setting in the CLI using the `config user fssso` command. See [Examples and troubleshooting on page 332](#).
5. Set **Collector Agent AD access mode** to **Advanced**, and select the **LDAP Server** (in this example, ADserver) you configured previously. See [Examples and troubleshooting on page 332](#).
6. Select the **Users** or **Groups** or **Organizational Units** tab to select the users, groups, OU that you want to monitor.
7. Select **OK**.

To specify the collector agent for FSSO - CLI:

```
config user fssso
  edit "WinGroups"
    set ldap-server "ADserver"
    set password ENC
      G7GQV7NEqilCM9jKmVmJJFVvhQ2+wtNEe9T0iYA5Sa+EqT2J8zhOrbkJFDr0RmY3c4LaoXdsoBczA
      ldONmcGfthTxxwGsigzGpbJdC7lspFlQYtj
    set server "10.11.101.160"
  end
```

To create the FSSO_Internet-users user group - web-based manager:

1. Go to **User & Device > User Groups** and select **Create New**.
2. Enter the following information and then select **OK**:

Name	FSSO_Internet_users
Type	Fortinet Single Sign-On (FSSO)
Members	Engineering, Sales

To create the FSSO_Internet-users user group - CLI:

```
config user group
  edit FSSO_Internet_users
    set group-type fsso-service
    set member CN=Engineering,cn=users,dc=office,dc=example,dc=com
      CN=Sales,cn=users,dc=office,dc=example,dc=com
  end
```

Creating the firewall user group

The non-FSSO users need a user group too. In this example, only two users are shown, but additional members can be added easily.

To create the firewall user group - web-based manager:

1. Go to **User & Device > User Groups** and select **Create New**.
2. Enter the following information and then select **OK**:

Name	Internet_users
Type	Firewall
Members	User1, User2

To create the firewall user group - CLI:

```
config user group
  edit Internet_users
    set group-type firewall
    set member User1 User2
  end
```

Defining policy addresses

1. Go to **Policy & Objects > Addresses**.
2. Create the following addresses:

Address Name	Internal_net
---------------------	--------------

Type	Subnet
Subnet / IP Range	10.11.102.0/24
Interface	Port 3

Address Name	Windows_net
Type	Subnet
Subnet / IP Range	10.11.101.0/24
Interface	Port 2

Creating security policies

Two security policies are needed: one for firewall group who connect through port3 and one for FSSO group who connect through port2.

To create a security policy for FSSO authentication - web-based manager:

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information:

Incoming Interface	Port2
Source Address	Windows_net
Source User(s)	FSSO_Internet_users
Outgoing Interface	Port1
Destination Address	all
Schedule	always
Service	ALL
NAT	ON
Security Profiles	Optionally, enable security profiles.

3. Select OK.

To create a security policy for FSSO authentication - CLI:

```
config firewall policy
edit 0
set srcintf port2
set dstintf port1
set srcaddr Windows_net
set dstaddr all
set action accept
```



```

set groups FSSO_Internet_users
set schedule always
set service ANY
set nat enable
end

```

To create a security policy for local user authentication - web-based manager:

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information:

Incoming Interface	Port3
Source Address	Internal_net
Source User(s)	Internet_users
Outgoing Interface	Port1
Destination Address	all
Schedule	always
Service	ALL
NAT	ON
Security Profiles	Optionally, enable security profiles.

3. Select **OK**.

To create a security policy for local user authentication - CLI:

```

config firewall policy
edit 0
set srcintf port3
set dstintf port1
set srcaddr internal_net
set dstaddr all
set action accept
set schedule always
set groups Internet_users
set service ANY
set nat enable
end

```

LDAP dial-in using member-attribute example

In this example, users defined in MicroSoft Windows Active Directory (AD) are allowed to set up a VPN connection simply based on an attribute that is set to TRUE, instead of based on their user group. In AD the "Allow Dialin" property is activated in the user properties, and this sets the `msNPAllowDialin` attribute to "TRUE".

This same procedure can be used for other member attributes, as your system requires.

To accomplish this with a FortiGate unit, member-attribute must be set. This can only be accomplished through the CLI - the option is not available through the web-based manager.

Before configuring the FortiGate unit, ensure the AD server has the `msNPAllowDialin` attribute set to "TRUE" for the users in question. If not, those users will not be able to authenticate.

To configure user LDAP member-attribute settings - CLI:

```
config user ldap
  edit "ldap_server"
    set server "192.168.201.3"
    set cnid "sAMAccountName"
    set dn "DC=fortilabanz,DC=com,DC=au"
    set type regular
    set username "fortigate@sample.com"
    set password *****
    set member-attr "msNPAllowDialin"
  next
end
```

To configure LDAP group settings - CLI:

```
config user group
  edit "ldap_grp"
    set member "ldap"
    config match
      edit 1
        set server-name "ldap"
        set group-name "TRUE"
      next
    end
  next
end
```

Once these settings are in place, users that are a member of the `ldap` user group will be able to authenticate.

To ensure your settings are correct, here is the sample output from a `diag debug` command that shows the authentication process.

When the "Allow Dial-in" attribute is set to "TRUE" the following will likely be in the output:

```
get_member_of_groups-Get the memberOf groups.
get_member_of_groups- attr='msNPAllowDialin', found 1 values
get_member_of_groups-val[0]='TRUE'
fnbamd_ldap_get_result-Auth accepted
fnbamd_ldap_get_result-Going to DONE state res=0
fnbamd_auth_poll_ldap-Result for ldap svr 192.168.201.3 is SUCCESS
fnbamd_auth_poll_ldap-Passed group matching
```

If the attribute is not set but it is expected, the following will likely be in the output:

```
get_member_of_groups-Get the memberOf groups.
get_member_of_groups- attr='msNPAllowDialin', found 1 values
get_member_of_groups-val[0]='FALSE'
fnbamd_ldap_get_result-Auth accepted
fnbamd_ldap_get_result-Going to DONE state res=0
fnbamd_auth_poll_ldap-Result for ldap svr 192.168.201.3 is SUCCESS
```

```
fnbamd_auth_poll_ldap-Failed group matching
```

The only difference between these two outputs is the last line which is either passed or failed based on if the member-attribute is set to the expected value or not.

RADIUS SSO example

A common RADIUS SSO topology involves a medium sized company network of users connecting to the Internet through the FortiGate unit, and authenticating with a RADIUS server. RADIUS SSO authentication was selected because it is fast and relatively easy to configure.

This section includes:

- [Assumptions](#)
- [Topology](#)
- [Configuring RADIUS](#)
- [Configuring FortiGate regular and RADIUS SSO security policies](#)
- [Testing](#)

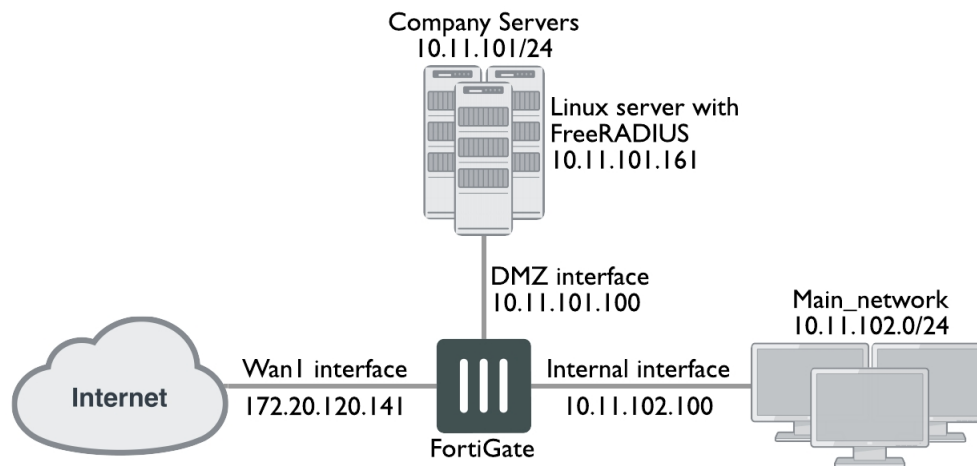
Assumptions

- VDOMs are not enabled.
- The admin super_admin administrator account will be used for all FortiGate unit configuration.
- Any other devices on the network do not affect the topology of this example, and therefore are not included.
- Anywhere settings are not described, they are assumed to be default values.
- A RADIUS server is installed on a server or FortiAuthenticator unit and uses default attributes.
- BGP is used for any dynamic routing.
- Authentication event logging under **Log & Report** has been configured.

Topology

Example.com has an office with 20 users on the internal network. These users need access to the Internet to do their jobs. The office network is protected by a FortiGate-60C unit with access to the Internet through the wan1 interface, the user network on the internal interface, and all the servers are on the DMZ interface. This includes an Ubuntu Linux server running FreeRADIUS. For this example only two users will be configured — Pat Lee with an account name `plee`, or `plee@example.com`, and Kelly Green with an account name `kgreen`, or `kgreen@example.com`.

RADIUS SSO topology



Configuring RADIUS

Configuring RADIUS includes configuring the RADIUS server such as FreeRADIUS, a radius client on user's computers, and configuring users in the system. For this example the two users will be Pat Lee, and Kelly Green. They belong to a group called `exampledotcom_employees`. When it is all configured, the RADIUS daemon needs to be started.

The users have a RADIUS client installed on their PCs that allows them to authenticate through the RADIUS server.

FreeRADIUS can be found on the freeradius.org website. For any problems installing FreeRADIUS, see the FreeRADIUS documentation.

Configuring FortiGate interfaces

Before configuring the RADIUS SSO security policy, configure FortiGate interfaces. This includes defining a DHCP server for the internal network as this type of network typically uses DHCP. The wan1 and dmz interfaces are assigned static IP addresses and do not need a DHCP server.

FortiGate interfaces used in this example

Interface	Subnet	Act as DHCP Server	Devices
wan1	172.20.120.141	No	Internet Service Provider
dmz	10.11.101.100	No	Servers, including RADIUS server
internal	10.11.102.100	Yes: x.x.x.110-.250	Internal user network

To configure FortiGate interfaces - web-based manager:

1. Go to **Network > Interfaces**.
2. Select wan1 to edit.

3. Enter the following information and select **OK**.

Alias	Internet
Addressing Mode	Manual
IP/Network Mask	172.20.120.141/255.255.255.0
Administrative Access	HTTPS, SSH
Enable DHCP Server	Not selected
Comments	Internet
Administrative Status	Up

4. Select dmz to edit.
5. Enter the following information and select **OK**.

Alias	Servers
Addressing Mode	Manual
IP/Network Mask	10.11.101.100/255.255.255.0
Administrative Access	HTTPS, SSH, PING, SNMP
Enable DHCP Server	Not selected
Listen for RADIUS Accounting Messages	Select
Comments	Servers
Administrative Status	Up

6. Select internal to edit.
7. Enter the following information and select **OK**.

Alias	Internal network
Addressing Mode	Manual
IP/Network Mask	10.11.102.100/255.255.255.0
Administrative Access	HTTPS, SSH, PING
Enable DHCP Server	Select
Address Range	10.11.102.110 - 10.11.102.250
Netmask	255.255.255.0

Default Gateway	Same as Interface IP
DNS Server	Same as System DNS
Comments	Internal network
Administrative Status	Up

Configuring a RADIUS SSO agent on the FortiGate unit

To create a RADIUS SSO agent:

1. Go to **Security Fabric > Fabric Connectors** and select **Create New**.
2. Under **SSO/Identity**, select **RADIUS Single Sign-On Agent**.
3. Enter a name for the RSSO Agent.
4. Enable **Use RADIUS Shared Secret** and enter the RADIUS server's shared secret.
5. Enable **Send RADIUS Responses**.
6. Select **OK**.

Creating a RADIUS SSO user group

To define a local user group for RADIUS SSO:

1. Go to **User & Device > User Groups** and select **Create New**.
2. Enter a Name for the user group.
3. In **Type**, select **RADIUS Single Sign-On (RSSO)**.
4. In **RADIUS Attribute Value**, enter the name of the RADIUS user group this local user group represents.
5. Select **OK**.

Configuring FortiGate regular and RADIUS SSO security policies

With the RADIUS server and FortiGate interfaces configured, security policies can be configured. This includes both RADIUS SSO and regular policies, as well as addresses and address groups. All policies require NAT to be enabled.

Security policies required for RADIUS SSO

Seq. No.	From -> To	Type	Schedule	Description
1	internal -> wan1	RADIUS SSO	business hours	Authenticate outgoing user traffic.
2	internal -> wan1	regular	always	Allow essential network services and VoIP.
3	dmz -> wan1	regular	always	Allow servers to access Internet.

Seq. No.	From -> To	Type	Schedule	Description
4	internal -> dmz	regular	always	Allow users to access servers.
5	any -> any	deny	always	Implicit policy denying all traffic that hasn't been matched.



The RADIUS SSO policy must be placed at the top of the policy list so it is matched first. The only exception to this is if you have a policy to deny access to a list of banned users. In this case, that policy must go at the top so the RADIUS SSO does not mistakenly match a banned user or IP address.

This section includes:

- [Schedules, address groups, and services groups](#)
- [Configuring regular security policies](#)
- [Configuring RADIUS SSO security policy](#)

Schedules, address groups, and services groups

This section lists the lists that need to be configured before security policies are created. Creating these lists is straight forward, so the essential information has been provided here but not step by step instructions. For more information on firewall related details, see

Schedules

Only one schedule needs to be configured — `business_hours`. This is a fairly standard Monday to Friday 8am to 5pm schedule, or whatever days and hours covers standard work hours at the company.

Address groups

The following address groups need to be configured before the security policies.

Address Group Name	Interface	Address range included
internal_network	internal	10.11.102.110 to 10.11.102.250
company_servers	dmz	10.11.101.110 to 10.11.101.250

Service groups

The following service groups need to be configured before the security policies. Note that the services listed are suggestions and may include more or less as required.

Service Group Name	Interface	Description of services to be included
essential_network_services	internal	Any network protocols required for normal network operation such as DNS, NTP, BGP.
essential_server_services	dmz	All the protocols required by the company servers such as BGP, HTTP, HTTPS, FTP, IMAP, POP3, SMTP, IKE, SQL, MYSQL, NTP, TRACEROUTE, SOCKs, and SNMP.
user_services	internal	Any protocols required by users HTTP, HTTP, FTP,

The following security policy configurations are basic and only include logging, and default AV and IPS.

Configuring regular security policies

Regular security policies allow or deny access for non-RADIUS SSO traffic. This is essential as there are network services—such as DNS, NTP, and FortiGuard—that require access to the Internet.

To configure regular security policies - web-based manager:

1. Go to **Policy & Objects > IPv4 Policy**, and select **Create New**.
2. Enter the following information, and select **OK**.

Incoming Interface	Internal
Source Address	internal_network
Outgoing Interface	wan1
Destination Address	all
Schedule	always
Service	essential_network_services
Action	ACCEPT
NAT	ON
Security Profiles	ON: AntiVirus, IPS
Log Allowed Traffic	ON
Comments	Essential network services

3. Select **Create New**, enter the following information, and select **OK**.

Incoming Interface	dmz
Source Address	company_servers
Outgoing Interface	wan1
Destination Address	all
Schedule	always
Service	essential_server_services
Action	ACCEPT
NAT	ON
Security Profiles	ON: AntiVirus, IPS
Log Allowed Traffic	enable
Comments	Company servers accessing the Internet

4. Select **Create New**, enter the following information, and select **OK**.

Incoming Interface	Internal
Source Address	internal_network
Outgoing Interface	dmz
Destination Address	company_servers
Schedule	always
Service	all
Action	ACCEPT
NAT	ON
Security Profiles	ON: AntiVirus, IPS
Log Allowed Traffic	enable
Comments	Access company servers

Configuring RADIUS SSO security policy

The RADIUS SSO policy allows access for members of specific RADIUS groups.

To configure RADIUS SSO security policy:

1. Go to **Policy & Objects > IPv4 Policy**.
2. Select **Create New**.
3. Enter the following information:

Incoming Interface	Internal
Source Address	internal_network
Source User(s)	Select the user groups you created for RSSO.
Outgoing Interface	wan1
Destination Address	all
Schedule	business_hours
Service	ALL
Action	ACCEPT
NAT	ON
Security Profiles	ON: AntiVirus, Web Filter, IPS, and Email Filter. In each case, select the default profile.

4. Select **OK**.
5. To ensure an RSSO-related policy is matched first, the policy should be placed higher in the security policy list than more general policies for the same interfaces.
6. Select **OK**.

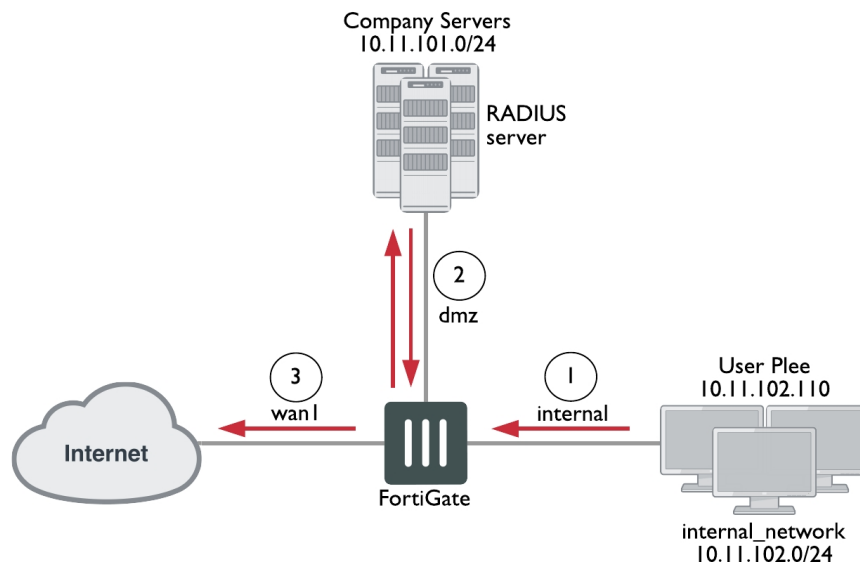
Testing

Once configured, a user only needs to log on to their PC using their RADIUS account. After that when they attempt to access an Internet website, the FortiGate unit will use their session information to get their RADIUS information. Once the user is verified, they are allowed access to the website.

To test the configuration perform the following steps:

1. Have user 'plee' logon to their PC, and try to access an Internet website.
2. The FortiGate unit will contact the RADIUS server for user plee's information.
Once confirmed, plee will have access to the website.
Each step generates log entries that enable you to verify that each step was successful.
3. If a step is unsuccessful, confirm that your configuration is correct.

RADIUS SSO test



Troubleshooting

In the web-based manager, a good tool for troubleshooting is the packet counter column on the security policy page at **Policy & Objects > IPv4 Policy**. This column displays the number of packets that have passed through this security policy. Its value when you are troubleshooting is that when you are testing your configuration (end to end connectivity, user authentication, policy use) watching the packet count for an increase confirms any other methods you may be using for troubleshooting. It provides the key of which policy is allowing the traffic, useful information if you expect a user to require authentication and it never happens.

This section addresses how to get more information from the CLI about users and user authentication attempts to help troubleshoot failed authentication attempts.

```
diag firewall iprope list
```

Shows the IP that the computer connected from. This is useful to confirm authorization and VPN settings.

```
diag firewall iprope clear
```

Clear all authorized users from the current list. Useful to force users to re-authenticate after system or group changes. However, this command may easily result in many users having to re-authenticate, so use carefully.

```
diag rso query ip
```

```
diag rso query rso-key
```

Queries the RSSO database.

For more information on troubleshooting specific features, go to that section of this document. Most sections have troubleshooting information at the end of the section. In addition to that information, see the [FortiOS Handbook Troubleshooting](#) guide for general troubleshooting information.

Chapter 4 - Best Practices

Introduction

This FortiGate Best Practices document is a collection of guidelines to ensure the most secure and reliable operation of FortiGate units in a customer environment. It is updated periodically as new issues are identified.

General considerations

1. For security purposes, NAT mode is preferred because all of the internal or DMZ networks can have secure private addresses. NAT mode policies use network address translation to hide the addresses in a more secure zone from users in a less secure zone.
2. Use virtual domains (VDOMs) to group related interfaces or VLAN subinterfaces. Using VDOMs will partition networks and create added security by limiting the scope of threats.
3. Use Transparent mode when a network is complex and does not allow for changes in the IP addressing scheme.

Customer service and technical support

For antivirus and IPS updates, firmware updates, updated product documentation, technical support information, and other resources, visit the Fortinet Support website at <https://support.fortinet.com>.

You can also register Fortinet products and service contracts and change your registration information at any time. You can also find a guide to Fortinet's Support services from: [How to work with Fortinet Support](#).

When requesting technical support, for optimum results you should provide as much of the following information as possible:

- Your name, and your company's name and location
- Your email address and/or telephone number
- Your support contract number (if applicable)
- The product name and model number
- The product serial number (if applicable)
- The software or firmware version number
- A detailed description of the problem

Fortinet Knowledge Base

The most recent Fortinet technical documentation is available from the Fortinet Knowledge Base. The knowledge base contains short how-to articles, FAQs, technical notes, product and feature guides, and much more. Visit the Fortinet Knowledge Base at <http://kb.fortinet.com>.

Comments on Fortinet technical documentation

Please send information about any errors or omissions in this document, or any Fortinet technical documentation, to techdoc@fortinet.com.

System and performance

By implementing the following best practices for system and performance, you will ensure maximum efficiency of your FortiGate device. Be sure to read everything carefully, particularly the section that concerns shutting down the FortiGate system, in order to avoid potential hardware issues.

Performance

- Disable any management features you do not need. If you don't need SSH or SNMP, disable them. SSH also provides another possibility for would-be hackers to infiltrate your FortiGate unit.
- Put the most used firewall rules to the top of the interface list.
- Log only necessary traffic. The writing of logs, especially if to an internal hard disk, slows down performance.
- Enable only the required application inspections.
- Keep alert systems to a minimum. If you send logs to a syslog server, you may not need SNMP or email alerts, making for redundant processing.
- Establish scheduled FortiGuard updates at a reasonable rate. Daily updates occurring every 4-5 hours are sufficient for most situations. In more heavy-traffic situations, schedule updates for the evening when more bandwidth can be available.
- Keep security profiles to a minimum. If you do not need a profile on a firewall rule, do not include it.
- Keep VDOMs to a minimum. On low-end FortiGate units, avoid using them if possible.
- Avoid traffic shaping if you need maximum performance. Traffic shaping, by definition, slows down traffic.

Shutting down

Always shut down the FortiGate operating system properly before turning off the power switch to avoid potentially catastrophic hardware problems.

To power off the FortiGate unit - web-based manager:

1. Go to **Dashboard**.
2. In the **System Resources** widget, select **Shutdown**.

To power off the FortiGate unit – CLI:

```
execute shutdown
```

Once this has been done, you can safely turn off the power switch or disconnect the power cables from the power supply.

Migration

Network administrators are often reluctant to change firewall vendors due to the perception that the migration process is difficult. Indeed, there is no point hiding the fact that moving to a new vendor requires careful consideration. But concern over the potential pain of migration should not stand in the way of adopting new security technologies. The purpose of this chapter is to describe the best practices for performing such migrations and ultimately to ease the migration process itself.

Information gathering

It is always best practice to perform a full network audit prior to any migration. This should include:

- Full back up of all security systems (including switches, routers) in case a back-out needs to be performed.
- Physical and logical network diagram with visual audit

Understanding exactly where cables run in the network and verifying they are all correctly labeled is essential to avoid mistakes and unnecessary downtime during the upgrade. Don't overlook simple things such as:

- Do I have enough spare interfaces on my switches?
- Do I have the right fiber (single/multi mode) and right connectors (LC, FC, MTRJ, SC, ST)?
- Do I have spare cables? (in the heat of the moment, it is a simple mistake to break an RJ-45 connector or damage a fiber)
- Do I have space in the rack for the new equipment?
- Do I have enough power sockets?

No matter how securely a FortiGate is configured in the network, it cannot help if it has been bypassed; visually checking where the device sits in the network in relation to other devices will ensure you are maintaining security and verify the network diagram is 'as built'. Details of all networks including subnet masks should be documented at this point to ensure that the replacement device is configured with the correct information.

Object and policy migration

Whilst we have suggested some level of manual review is included in the policy migration, it can be useful to be able to automatically migrate simply between another vendor's format and the FortiGate format. The FortiGate policy format is text based and can easily be cut and pasted into from other vendor formats however, responding to the high customer demand to migrate away from other vendors, Fortinet have released an automatic configuration migration tool at <http://convert.fortinet.com> to simplify this process. Supporting Cisco ACLs, PIX, ASA, Check Point, and Juniper, the Converter can securely upload and convert the policy into the Fortinet format.

Testing and validation

This is an important process and should be tested offline first wherever possible i.e. configure the policy in the lab or on a test network and verify that the required access permissions are being implemented. To really test the solution out, the FortiGate can be implemented on the live network with a different gateway IP and the selected user pointed to the new gateway. This allows a staged approach to migrating the new platform into the network ensuring that the process does not interrupt day to day operations.

Going live and obtaining feedback

If testing and validation is successful at this point, you can migrate to the new firewall either by switching IP's and removing the old devices or by changing the default gateway in DHCP. Once the firewall is in place, acceptance testing will of course need to be carried out and an iterative process of tuning undertaken to finalize the configuration.

Adding new services

The Fortinet solution will have a plethora of additional features compared to your previous vendor and it is very tempting to start switching them on but it is a good idea to wait and validate the new firewall as was previously configured before adding new functions as this simplifies testing and problem diagnosis. Finally complete the migration (don't forget about the Plan Do Check Act Cycle) by adding any new services that were requested and learn about the multiple features you have available with the FortiGate appliance.

Environmental specifications

Keep the following environmental specifications in mind when installing and setting up your FortiGate unit.

- Operating temperature: 32 to 104°F (0 to 40°C). Temperatures may vary, depending on the FortiGate model.
- If you install the FortiGate unit in a closed or multi-unit rack assembly, the operating ambient temperature of the rack environment may be greater than room ambient temperature.
Therefore, make sure to install the equipment in an environment compatible with the manufacturer's maximum rated ambient temperature.
- Storage temperature: -13 to 158°F (-25 to 70°C). Temperatures may vary, depending on the FortiGate model.
- Humidity: 5 to 90% non-condensing.
- Air flow - For rack installation, make sure that the amount of air flow required for safe operation of the equipment is not compromised.
- For free-standing installation, make sure that the appliance has at least 1.5 in. (3.75 cm) of clearance on each side to allow for adequate air flow and cooling.

Depending on your device, the FortiGate may generate, use, and even radiate radio frequency energy and, if not installed and used in accordance with the instructions, may cause harmful interference to radio communications. However, there is no guarantee that interference will not occur in a particular installation. If the equipment does cause harmful interference to radio or television reception, which can be determined by turning the equipment off and on, the user is encouraged to try to correct the interference by one or more of the following measures:

- Reorient or relocate the receiving antenna.
- Increase the separation between the equipment and receiver.
- Connect the equipment into an outlet on a circuit different from that to which the receiver is connected.
- Consult the dealer or an experienced radio/TV technician for help.



Explosion is a serious risk if the battery is replaced by an incorrect type. Dispose of used batteries according to the instructions. To reduce the risk of fire, use only No. 26 AWG or larger UL Listed or CSA Certified Telecommunication Line Cord.

Grounding

- Ensure the FortiGate unit is connected and properly grounded to a lightning and surge protector. WAN or LAN connections that enter the premises from outside the building should be connected to an Ethernet CAT5 (10/100 Mb/s) surge protector.
- Shielded Twisted Pair (STP) Ethernet cables should be used whenever possible rather than Unshielded Twisted Pair (UTP).
- Do not connect or disconnect cables during lightning activity to avoid damage to the FortiGate unit or personal injury.

Rack mounting

- **Elevated Operating Ambient** - If installed in a closed or multi-unit rack assembly, the operating ambient temperature of the rack environment may be greater than room ambient.
Therefore, consideration should be given to installing the equipment in an environment compatible with the maximum ambient temperature (Tmax) specified by the manufacturer.

- **Reduced Air Flow** - Installation of the equipment in a rack should be such that the amount of air flow required for safe operation of the equipment is not compromised.
- **Mechanical Loading** - Mounting of the equipment in the rack should be such that a hazardous condition is not achieved due to uneven mechanical loading.
- **Circuit Overloading** - Consideration should be given to the connection of the equipment to the supply circuit and the effect that overloading of the circuits might have on overcurrent protection and supply wiring. Appropriate consideration of equipment nameplate ratings should be used when addressing this concern.
- **Reliable Earthing** - Reliable earthing of rack-mounted equipment should be maintained.

Particular attention should be given to supply connections other than direct connections to the branch circuit (e.g. use of power strips).

Firmware

Firmware upgrading and downgrading sounds pretty simple, anyone can do it, right? The mark of a professional is not that they can do something correctly, or even do it correctly over and over again. A professional works in such a way that, if anything goes wrong they are prepared and able to quickly get things back to normal. Firmware updates can go wrong just like anything else. So a real professional does things in a way that minimizes their risk and follows some best practices, as listed below.

Firmware change management

Consider the following five points when performing firmware upgrades, not only in FortiOS but in general. This applies to pretty much any change you have to do in a production environment.

Understanding the new version first

Before attempting any changes in production, first make sure you set up a laboratory where you can freely play with the new features, and understand them with enough time and no pressure. Read the Release Notes, Manuals, and other documentation like presentations, videos, or podcasts about the new version.

You are ready to explain the need for an upgrade once you understand:

- The differences and the enhancements between the new version and the previous version(s).
- The impact of the upgrade on customers and the users of the operating platform.
- The known limitations that might affect your environment.
- The potential risks when performing the upgrade.
- The licensing changes that may apply.



Never attempt to upgrade to a version you don't fully understand (both on features and known limitations), and on which you have no operational experience.

Have a valid reason to upgrade

The reason can NOT be "Because I want to have the latest version". The reason has to be explained in terms of business, technical, and/or operational improvement.

Affirmative answers to the following questions are valid reasons to upgrade:

- Does the new version have a feature that helps to ensure compliance?
- Does the new version have an enhancement that allows 40% decrease (40% improvement) on the time to perform a certain operation?
- Does the new feature correct a known defect/bug found on a previous version that affects the company business/operations?
- Will the new version allow your organization to deploy new services that will help to gain new customers or increase loyalty of existing ones?
- Is the vendor cutting support for the version your organization is currently using?

If the best reason to upgrade is “Because the new features seem to be cool” or “Because I want to have the latest version”, a little more understanding and planning may be necessary.

Prepare an upgrade plan

If you choose to upgrade because you found a valid reason to do so, make sure you create a plan that covers business, technical, and operational aspects of the upgrade:

Business:

Proper planning and justification for an upgrade should be proportional to how critical the system is to the business.

- Make sure you can clearly articulate the benefits of the upgrade in business terms (time, money, and efficiency).
- Understand the business processes that will be affected by the change.
- Make sure the upgrade maintenance window is not close to a business-critical process (such as quarterly or monthly business closure).
- Obtain executive and operational approval for the maintenance window. The approval must come from the owners of ALL the systems/information affected by the upgrade, not only from those that own the system being upgraded. The approval must be done in a formal (written or e-mail) form.

Technical and operational:

- Re-read the Release Notes for the technology you are upgrading. Supported hardware models, upgrade paths, and known limitations should be clearly understood.
- Make sure your upgrade maintenance window does not overlap with any other maintenance window on your infrastructure.
- If you have any premium support offer (such as TAM, Premium Support), do a capacity planning exercise to ensure the new firmware/software version does not take more hardware resources than you currently have.
- Create a backup, whether or not you have scheduled backups. Create a new fresh backup.
- Obtain offline copies of both the currently installed firmware and the new version.
- Create a list of systems with inter-dependencies to the system you are upgrading. For example, if you are upgrading a FortiGate; understand the impact on any FortiAP, FortiAuthenticator, FortiToken, FortiManager, or FortiAnalyzer you have on your environment.
- Ensure you have a list of adjacent devices to the upgrading platform and have administrative access to them, just in case you need to do some troubleshooting. Are you upgrading FortiWeb? Make sure you can administratively access the Web Applications. Are you upgrading a FortiGate? Make sure you can administratively access the surrounding switches and routers.
- Have a step-by-step plan on how to perform and test the upgrade. You want to make sure you think of the worst situation before it happens, and have predefined courses of action, instead of thinking under pressure when something already went wrong.
- Define a set of tests (that include critical business applications that should be working) to make sure the upgrade went fine. If any test does not go well, define which ones mandate a rollback and which ones can be tolerated for further troubleshooting. This set of tests should be run before and after the upgrade to compare results, and they should be the same.
- Define a clear rollback plan. If something goes wrong with the upgrade or the tests, the rollback plan will help you get your environment back to a known and operational status. The plan must clearly state the conditions under which the rollback will be started.

- Declare configuration freezes. A little bit before and after the upgrade. The idea is to reduce the amount of variables to take into consideration if something goes wrong.
- Perform a “Quality Assurance” upgrade. Grab a copy of the production configuration, load it on a non-production box and execute the upgrade there to see if there are any issues on the process. Then adjust your plan according to the results you obtained.
- Have a list of information elements to be gathered if something goes wrong. This ensures that, even if the upgrade fails, you will collect enough information so you can troubleshoot the issue without needing to repeat the problem. Get help from Fortinet Support if you need to check what else could be missing on your list.
- Define a test monitoring period after the change was completed. Even if the upgrade went smoothly, something could still go wrong. Make sure you monitor the upgraded system for at least one business cycle. Business cycles may be a week, a month, or a quarter, depending on your organization’s business priorities.

Execute the upgrade plan

Execution of an upgrade is just as key as planning.

Once you are performing the upgrade, the pressure will rise and stress might peak. This is why you should stick to the plan you created with a cool head.

Resist the temptation to take decisions while performing the upgrade, as your judgment will be clouded by the stress of the moment, even if a new decision seems to be “obvious” at such time. If your plan says you should rollback, then execute the rollback despite the potential “We-can-fix-this-very-quickly” mentality.

While performing the upgrade, make sure all the involved components are permanently monitored before, during, and after the upgrade, either via monitoring systems, SNMP alerts, or at least with tools like a ping. Critical resources like CPU, memory, network, and/or disk utilization must also be constantly monitored.

To avoid misunderstandings, when performing the tests for each critical application defined on the planning, make sure there are formal notifications on the results for each user area, service, system, and/or application tested.

Regardless if you have to rollback or not, if a problem occurs, make sure you gather as much information about the problem as possible, so you can later place a Support ticket to find a solution.

Last but not least, document the upgrade:

- Enable your terminal emulation program to leave trace of all the commands executed and all the output generated. If you are performing steps via GUI, consider using a video capture tool to document it.
- Document any command or change performed over the adjacent/interdependent systems. Make sure they are acknowledged by the relevant administrators
- Document any deviations performed over the upgrade plan. This is planned-versus-actual.

Learn more about change management

Change Management and Change Control are huge knowledge areas in the field of Information Systems and Computer/Network Security.

This document is by no means a comprehensive list on what you should do when performing an upgrade, with either Fortinet or any other technology. It is merely a list of important things you should take into consideration when performing upgrades which are the result of years of experience dealing with changes on critical environments, as it is common that security devices are protecting critical applications and processes.

There are vast resources on the topic: books, public white papers, blog entries, etc. If you search the Internet for the “Change Control Best Practices” or “Change Management Best Practices” you will get many interesting documents.



Changes on production IT infrastructure are critical to the business. Make sure they play in your favor and not against you.

Performing a firmware upgrade

Upgrading a firewall is something that should be compared to upgrading the operating system on your computer. It's not to be taken lightly! You want to make sure everything is backed up and you have some options available if things go awry. Assuming it all seems to work you also want a list of things to do in order to confirm everything is working properly. Finally, you need enough time to do it. All really simple stuff, but what does this mean in relation to upgrading your FortiGate? It means, you follow these simple steps:

- 1. Backup and store old configuration (full configuration backup from CLI).**

Digging into this a little, step 1 is easy to understand. Do a full backup of your old configuration. This is all part of your disaster recovery plan. If the upgrade fails in some way you need to make sure you can get the Firewall back up and running. The best way to do this is to get it back to a state where you know what the behavior was. For more information, refer to ["Performing a configuration backup" on page 359](#).

- 2. Have copy of old firmware available.**

Step 2, is also part of your disaster recovery. If the upgrade fails you might be able to switch the active partition. But as a Professional, you need to be prepared for the worst case scenario where you can't do that. Which means you'll need your old firmware.

- 3. Have disaster recovery option on standby – especially if remote.**

Step 3, is your plan for what to do in the event of a critical failure. As we're talking FortiGate this means that your firewall doesn't come back after the upgrade. What this means is that you need to be able to get to the console port in order to find out why. Maybe it's DHCP and the IP changed, maybe the OS is corrupt, who knows? Get to the console and find out.

There could be a simple fix. If there's not, then be prepared for a format and TFTP reload.

- 4. Read the release notes, including the upgrade path and bug information.**

Step 4, READ THE RELEASE NOTES. They contain all kinds of information, known bugs, fixed bugs even upgrade issues like lost configuration settings. Not all upgrade information is ever contained in any products release notes. That does not mean they are devoid of good/useful information. Read them, digest them, then a few days later read them again.

- 5. Double check everything.**

Step 5, do a double check of everything. Is your TFTP server working, does your console connection function, is there anything in the release notes that could impact your upgrade procedure, do you have your configuration backed up? Make sure you've done everything.

- 6. Upgrade.**

Step 6, do the upgrade. Doing an upgrade doesn't take very long, a few minutes (less a lot of times) but make sure you schedule enough time for it. At the end of the day an upgrade can succeed or fail. If it succeeds you want some time to check/confirm that any important features you have are working (VPNs etc). If it fails you'll need time to sort things out.

Performing a firmware downgrade

Just like upgrading, you need to make sure it's done properly. While similar, the steps are somewhat different since there are other pitfalls in this case.

- 1. Locate pre-upgrade configuration file.**

Step 1 is very important. This is why, when you upgrade you make a backup of your old configuration and save it. If you don't, then you'll need to rebuild manually.

- 2. Have copy of old firmware available.**

Step 2 is fairly obvious. Even with devices that have multiple partitions and your downgrade process is simply going to be to switch the active partition, this could go wrong. In which case, you may be without Internet access. A professional has a plan for when things go wrong.

- 3. Have disaster recovery option on standby – especially if remote.**

Step 3 is no different from before. Hopefully you don't need to format the unit, but be prepared for that, just in case.

- 4. Read the release notes – is a downgrade possible, or necessary?**

Step 4, once again, is to READ THE RELEASE NOTES. In this case, you will need to do this for the version you are on, and the version you are downgrading too, and everything in between (if you are going back multiple major releases or patches). Maybe the OS switched from 32 to 64 bits somewhere between the two firmware releases. In order to make sure you don't get nailed by something like that you need to check the upgrade and downgrade information in every major release and patch, as it may have a direct impact on your options.

- 5. Double check everything.**

- 6. Downgrade – all settings, except those needed for access, are lost.**

Step 5 and 6 are the same as before. Double check everything, then downgrade.

- 7. Restore pre-upgrade configuration.**

Step 7 is new. Obviously most settings are lost when you downgrade so in order to get back up and running you will need to restore your old configuration file.

Performing a configuration backup

Once you configure the FortiGate unit and it is working correctly, it is extremely important that you backup the configuration. In some cases, you may need to reset the FortiGate unit to factory defaults or perform a TFTP upload of the firmware, which will erase the existing configuration. In these instances, the configuration on the device will have to be recreated, unless a backup can be used to restore it.

It is also recommended that once any further changes are made that you backup the configuration immediately, to ensure you have the most current configuration available. Also, ensure you backup the configuration before upgrading the FortiGate unit's firmware. Should anything happen during the upgrade that changes the configuration, you can easily restore the saved configuration.

Always backup the configuration and store it on the management computer or off-site. You have the option to save the configuration file to various locations including the local PC, USB key, FTP and TFTP site. The latter two are configurable through the CLI only.

If you have VDOMs, you can back up the configuration of the entire FortiGate unit or only a specific VDOM. Note that if you are using FortiManager or FortiCloud, full backups are performed and the option to backup individual VDOMs will not appear.

To back up the FortiGate configuration - web-based manager:

1. Go to **Dashboard**.
2. On the **System Information** widget, select **Backup** next to **System Configuration**.
3. Select to backup to your **Local PC** or to a **USB Disk**.
The **USB Disk** option will be grayed out if no USB drive is inserted in the USB port. You can also backup to the FortiManager using the CLI.
4. If VDOMs are enabled, select to backup the entire FortiGate configuration (**Full Config**) or only a specific VDOM configuration (**VDOM Config**).
5. If backing up a VDOM configuration, select the VDOM name from the list.
6. Select **Encrypt configuration file**.
Encryption must be enabled on the backup file to back up VPN certificates.
7. Enter a password and enter it again to confirm it. You will need this password to restore the file.
8. Select **Backup**.
9. The web browser will prompt you for a location to save the configuration file. The configuration file will have a .conf extension.

To back up the FortiGate configuration - CLI:

```
execute backup config management-station <comment>
```

... or ...

```
execute backup config usb <backup_filename> [<backup_password>]
```

... or for FTP (note that port number, username are optional depending on the FTP site)...

```
execute backup config ftp <backup_filename> <ftp_server> [<port>] [<user_name>]  
[<password>]
```

... or for TFTP ...

```
execute backup config tftp <backup_filename> <tftp_servers> <password>
```

Use the same commands to backup a VDOM configuration by first entering the commands:

```
config vdom  
edit <vdom_name>
```

Backing up a configuration file using SCP

You can use secure copy protocol (SCP) to download the configuration file from the FortiGate unit as an alternative method of backing up the configuration file or an individual VDOM configuration file. This is done by enabling SCP for an administrator account and enabling SSH on a port used by the SCP client application to connect to the FortiGate unit. SCP is enabled using the CLI commands:

```
config system global  
set admin-scp enable  
end
```

Use the same commands to backup a VDOM configuration by first entering the commands:

```
config global  
set admin-scp enable  
end
```

```
config vdom
  edit <vdom_name>
```


Security Profiles (AV, Web Filtering etc.)

Infection can come from many sources and have many different effects. Because of this, there is no single means to effectively protect your network. Instead, you can best protect your network with the various UTM tools your FortiGate unit offers.

Firewall

- Be careful when disabling or deleting firewall settings. Changes that you make to the firewall configuration using the GUI or CLI are saved and activated immediately.
- Arrange firewall policies in the policy list from more specific to more general. The firewall searches for a matching policy starting from the top of the policy list and working down. For example, a very general policy matches all connection attempts. When you create exceptions to a general policy, you must add them to the policy list above the general policy.
- Avoid using the All selection for the source and destination addresses. Use addresses or address groups.
- If you remove all policies from the firewall, there are no policy matches and all connections are dropped.
- If possible, avoid port ranges on services for security reasons.
- The settings for a firewall policy should be as specific as possible. Do not use 0.0.0.0 as an address. Do not use Any as a service. Use subnets or specific IP addresses for source and destination addresses and use individual services or service groups.
- Use a 32-bit subnet mask when creating a single host address (for example, 255.255.255.255).
- Use logging on a policy only when necessary and be aware of the performance impact. For example, you may want to log all dropped connections but can choose to use this sparingly by sampling traffic data rather than have it continually storing log information you may not use.
- It is possible to use security policies based on 'any' interface. However, for better granularity and stricter security, explicit interfaces are recommended.
- Use the comment field to input management data, for example: who requested the rule, who authorized it, etc.
- Avoid FQDN addresses if possible, unless they are internal. It can cause a performance impact on DNS queries and security impact from DNS spoofing.
- For non vlan interfaces, use zones (even if you have only one single interface for members) to allow:
 - An explicit name of the interface to use in security policies ('internal' is more explicit than 'port10').
 - A split between the physical port and its function to allow port remapping (for instance moving from a 1G interface to a 10G interface) or to facilitate configuration translation, as performed during hardware upgrades.

Security

- Use NTP to synchronize time on the FortiGate and the core network systems, such as email servers, web servers, and logging services.
- Enable log rules to match corporate policy. For example, log administration authentication events and access to systems from untrusted interfaces.
- Minimize adhoc changes to live systems, if possible, to minimize interruptions to the network. When not possible, create backup configurations and implement sound audit systems using FortiAnalyzer and FortiManager.
- If you only need to allow access to a system on a specific port, limit the access by creating the strictest rule possible.

Authentication

- You must add a valid user group to activate the Authentication check box on the firewall policy configuration page.
- Users can authenticate with the firewall using HTTP or FTP. For users to be able to authenticate, you must add an HTTP or FTP policy that is configured for authentication.

Antivirus

- Enable antivirus scanning at the network edge for all services.
- Use FortiClient endpoint antivirus scanning for protection against threats that get into your network.
- Subscribe to FortiGuard AntiVirus Updates and configure your FortiGate unit to receive push updates. This will ensure you receive antivirus signature updates as soon as they are available.
- To ensure that all AV push updates occur, ensure you have an AV profile enabled in a security policy.
- Enable only the protocols you need to scan. If you have antivirus scans occurring on the SMTP server, or use FortiMail, it is redundant to have scanning occur on the FortiGate unit as well.
- Reduce the maximum file size to be scanned. Viruses usually travel in small files of around 1 to 2 megabytes.
- Do not quarantine files unless you regularly monitor and review them. This is otherwise a waste of space and impacts performance.
- Examine antivirus reports and log messages periodically. Take particular notice of repeated detections. For example, repeated virus detection in SMTP traffic could indicate a system on your network is infected and is attempting to contact other systems to spread the infection using a mass mailer.

Antispam

- If possible use, a FortiMail unit. The antispam engines are more robust.
- Use fast DNS servers.
- Use specific security profiles for the rule that will use antispam.
- DNS checks may cause false positive with HELO DNS lookup.
- Content analysis (banned words) may impose performance overhead.

Intrusion Prevention System (IPS)

Your FortiGate's IPS system can detect traffic attempting to exploit this vulnerability. IPS may also detect when infected systems communicate with servers to receive instructions. Refer to the following list of best practices regarding IPS.

- Enable IPS scanning at the network edge for all services.
- Use FortiClient endpoint IPS scanning for protection against threats that get into your network.
- Subscribe to FortiGuard IPS Updates and configure your FortiGate unit to receive push updates. This will ensure you receive IPS signature updates as soon as they are available.
- Because it is critical to guard against attacks on services that you make available to the public, configure IPS signatures to block matching signatures. For example, if you have a web server, configure the action of web server signatures to Block.
- Create and use security profiles with specific signatures and anomalies you need per-interface and per-rule.
- Do not use predefined or generic profiles. While these profiles are convenient to supply immediate protection, you should create profiles to suit your network environment.

- If you do use the default profiles, reduce the IPS signatures/anomalies enabled in the profile to conserve processing time and memory.
- If you are going to enable anomalies, make sure you tune thresholds according to your environment.
- If you need protection, but not audit information, disable the logging option.
- Tune the IP-protocol parameter accordingly.

Blocking Skype using CLI options for improved detection

If you want to identify or block Skype sessions, use the following CLI command with your FortiGate's public IP address to improve detection (FortiOS 4.3.12+ and 5.0.2+):

```
config ips global
    set skype-client-public-ipaddr 198.51.100.0,203.0.113.0
end
```

Note that the above syntax is configured using multiple public IP addresses, where a single public IP address may suffice depending on your network configuration.

Email filter

Spam is a common means by which attacks are delivered. Users often open email attachments they should not, and infect their own machine.

- Enable email filtering at the network edge for all types of email traffic.
- Use FortiClient endpoint IPS scanning for protection against threats that get into your network.
- Subscribe to the FortiGuard AntiSpam Service.

URL filtering

Best practices for URL filtering can be divided into four categories: flow-based versus proxy based filtering; local category/rating feature; URL filter 'Exempt' action; and Deep Scan.

Flow-based versus proxy-based

Try to avoid mixing flow-based and proxy-based features in the same profile if you are not using IPS or Application Control.

Local category/rating feature

Local categories and local rating features consume a large amount of CPU resources, so use these features as little as possible. It is better to use Local categories instead of using the 'override' feature, since the 'override' feature is more complicated and more difficult to troubleshoot.

URL filter 'Exempt' action

When using the URL filter 'Exempt' option, webfilter, antivirus and dlp scans are bypassed by default, so use this option only for trusted sites.

Configuration notes: You need to configure 'Exempt' actions in the URL filter if you want to bypass the FortiGuard Web Filter. You can configure which particular inspection(s) you want to bypass using the `set exempt` command in `config webfilter urlfilter`.

Deep Scan

The 'Deep Scan' feature is much heavier on resources than 'HTTPS URL Scan Only'. Deep Scan is much more accurate, since many sites (such as various Google applications) cannot be scanned separately without deep scanning enabled.

Note: If you configure Deep Scan in the SSL profile and then configure 'Enable HTTPS URL Scan Only' in the web filter profile, then Deep Scan is not performed.

Web filtering

FortiGuard Web Filtering can help stop infections from malware sites and help prevent communication if an infection occurs.

- Enable FortiGuard Web Filtering at the network edge.
- Install the FortiClient application and use FortiGuard Web Filtering on any systems that bypass your FortiGate unit.
- Block categories such as Pornography, Malware, Spyware, and Phishing. These categories are more likely to be dangerous

Patch management

When vulnerabilities are discovered in software, the software vendors release updates that fix these problems. Keeping your software and operating system up-to-date is a vital step to prevent infection and defend against attacks.

- Follow the latest advisories and reports on the FortiGuard webpage.
- Apply updates to all software as the updates become available.
- FortiGuard Vulnerability Management can help identify security weaknesses in your network. This subscription service is available through FortiScan and FortiAnalyzer units.
- Apply firmware updates to your FortiGate unit as they are released.
- Subscribe to FortiGuard AntiVirus and IPS services, so that AntiVirus and IPS scanning engines are automatically updated when new version are released.

Policy configuration

Configuring the FortiGate unit with an 'allow all' traffic policy is very undesirable. While this does greatly simplify the configuration, it is less secure. As a security measure, it is best practice for the policy rulebase to 'deny' by default, and not the other way around.

Policy configuration changes

On a heavy-loaded system, plan configuration changes during low usage periods in order to minimize impact on CPU usage and established sessions. In this scenario, it is considered a best practice to de-accelerate the hardware-accelerated sessions.

You can configure de-accelerated behaviour on hardware-accelerated sessions using CLI commands to control how the processor manages policy configuration changes. The following CLI commands are to be used:

```
config system settings
    set firewall-session-dirty { check-all | check-new | check-policy-option }
end
```

where you want the following to be true:

<code>check-all</code>	CPU flushes all current sessions and re-evaluates them. This is the default option.
<code>check-new</code>	CPU keeps existing sessions and applies policy changes to new sessions only. This reduces CPU load and the possibility of packet loss.
<code>check-policy-option</code>	Use the option selected in the <code>firewall-session-dirty</code> field of the firewall policy (<code>check-all</code> or <code>check-new</code> , as above, but per policy).

Policy whitelisting

- Allow only the necessary inbound and outbound traffic.
- If possible, limit traffic to specific addresses or subnets. This allows the FortiGate unit to drop traffic to and from unexpected addresses.

IPS and DoS policies

- Because it is critical to guard against attacks on services that you make available to the public, configure IPS signatures to block matching signatures. For example, if you have a web server, configure the action of web server signatures to Block.
- Your FortiGate unit includes IPS signatures written to protect specific software titles from DoS attacks. Enable the signatures for the software you have installed and set the signature action to Block.
- DoS attacks are launched against vulnerabilities. Maintain a FortiGuard IPS subscription to ensure your FortiGate unit automatically receives new and updated IPS signatures as they are released.
- Use and configure DoS policies to appropriate levels based on your network traffic and topology. This will help drop traffic if an abnormal amount is received. The key is to set a good threshold. The threshold defines the maximum number of sessions/packets per second of normal traffic. If the threshold is exceeded, the action is triggered. Threshold defaults are general recommendations, but your network may require very different values. One way to find the correct values for your environment is to set the action to Pass and enable logging. Observe the logs and adjust the threshold values until you can determine the value at which normal traffic begins to generate attack reports. Set the threshold above this value with the margin you want. Note that the smaller the margin, the more protected your system will be from DoS attacks, but your system will also be more likely to generate false alarms.

Networking

When configuring your network, ensure that there is no 'back door' access to the protected network. For example, if there is a wireless access point, it must be appropriately protected with password and encryption.

Be sure to also maintain an up-to-date network diagram which includes IP addressing, cabling, and network elements.

Routing configuration

- Always configure a default route.
- Add blackhole routes for subnets reachable using VPN tunnels. This ensures that if a VPN tunnel goes down, traffic is not mistakenly routed to the Internet unencrypted.

Policy routing

Keep the number of policy routes to a minimum to optimize performance in route lookup and to simplify troubleshooting.

Dynamic routing

- Select a Router ID that matches an IP assigned to an interface. This avoids the likelihood of having two devices with the same router ID.
- For routing over an IPsec tunnel, assign IP addresses to both ends of the tunnel.

Advanced routing

Use the following best practices for advanced routing when dealing with Border Gateway Protocol (BGP) and Open Shortest Path First (OSPF).

Border Gateway Protocol (BGP)

If you are using BGP, it is recommended that you enable soft-reconfiguration. This has two benefits:

- It allows you to perform 'soft clear' of peers after a change is made to a BGP policy.
- It provides greater visibility into the specific prefixes learned from each neighbor.

Leave soft-reconfiguration disabled if your FortiGate does not have much unused memory. Soft-reconfiguration requires keeping separate copies of prefixes received and advertised, in addition to the local BGP database.

Open Shortest Path First (OSPF)

- Avoid use of passive interfaces wherever possible.
- Avoid use of virtual links to connect areas. All areas should be designed to connect directly to the backbone area.
- Ensure that all backbone routers have a minimum of two peering connections to other backbone neighbors.
- An entire OSPF domain should be under common administration.

Network Address Translation (NAT)

- Beware of misconfiguring the IP Pool range. Double-check the start and end IPs of each IP pool. The IP pool should not overlap with addresses assigned to FortiGate interfaces or to any hosts on directly connected networks.
- If you have internal and external users accessing the same servers, use split DNS to offer an internal IP to internal users so that they don't have to use the external-facing VIP.

Configuring NAT

Do not enable NAT for inbound traffic unless it is required by an application. If, for example, NAT is enabled for inbound SMTP traffic, the SMTP server might act as an open relay.

Transparent Mode

- Do not connect two ports to the same VLAN on a switch or to the same hub. Some Layer 2 switches become unstable when they detect the same MAC address originating on more than one switch interface or from more than one VLAN.
- If you operate multiple VLANs on your FortiGate unit, assign each VLAN id to its own forwarding domain to ensure that the scope of the broadcast does not extend beyond the VLAN it originated in.

To protect against Layer 2 loops:

- Enable `stpforward` on all interfaces.
- Use separate VDOMs for production traffic (TP mode VDOM) and management traffic (NAT/Route mode VDOM).
- Only place those interfaces used for production in the TP mode VDOM. Place all other interfaces in the NAT/Route mode VDOM. This protects against potential Layer 2 loops.

Using virtual IPs (VIPs)

- Use the external IP of 0.0.0.0 when creating a VIP for a FortiGate unit where the external interface IP address is dynamically assigned.
- Be sure to select the correct external interface when creating a new virtual IP (VIP). The external interface should be set to the interface at which the FortiGate unit receives connection requests from external networks.

FGCP high availability

Fortinet suggests the following practices related to high availability:

- Use Active-Active HA to distribute TCP and UTM sessions among multiple cluster units. An active-active cluster may have higher throughput than a standalone FortiGate unit or than an active-passive cluster.
- Use a different host name on each FortiGate unit when configuring an HA cluster. Fewer steps are required to add host names to each cluster unit before configuring HA and forming a cluster.
- Consider adding an Alias to the interfaces used for the HA heartbeat so that you always get a reminder about what these interfaces are being used for.
- Enabling `load-balance-all` can increase device and network load since more traffic is load-balanced. This may be appropriate for use in a deployment using the firewall capabilities of the FortiGate unit and IPS but no other content inspection.
- An advantage of using session pickup is that non-content inspection sessions will be picked up by the new primary unit after a failover. The disadvantage is that the cluster generates more heartbeat traffic to support session pickup as a larger portion of the session table must be synchronized. Session pickup should be configured only when required and is not recommended for use with SOHO FortiGate models. Session pickup should only be used if the primary heartbeat link is dedicated (otherwise the additional HA heartbeat traffic could affect network performance).
- If session pickup is not selected, after a device or link failover all sessions are briefly interrupted and must be re-established at the application level after the cluster renegotiates. For example, after a failover, users browsing the web can just refresh their browsers to resume browsing. Users downloading large files may have to restart their download after a failover. Other protocols may experience data loss and some protocols may require sessions to be manually restarted. For example, a user downloading files with FTP may have to either restart downloads or restart their FTP client.
- If you need to enable session pickup, consider enabling `session-pickup-delay` to improve performance by reducing the number of sessions that are synchronized. See [Improving session synchronization performance on page 1](#).
- Consider using the `session-sync-dev` option to move session synchronization traffic off the HA heartbeat link to one or more dedicated session synchronization interfaces. See [Improving session synchronization performance on page 1](#).
- To avoid unpredictable results, when you connect a switch to multiple redundant or aggregate interfaces in an active-passive cluster you should configure separate redundant or aggregate interfaces on the switch; one for each cluster unit.
- Use SNMP, syslog, or email alerts to monitor a cluster for failover messages. Alert messages about cluster failovers may help find and diagnose network problems quickly and efficiently.

Heartbeat interfaces

Fortinet suggests the following practices related to heartbeat interfaces:



Do not use a FortiGate switch port for the HA heartbeat traffic. This configuration is not supported.

- Configure at least two heartbeat interfaces and set these interfaces to have different priorities.
- For clusters of two FortiGate units, as much as possible, heartbeat interfaces should be directly connected using patch cables (without involving other network equipment such as switches). If switches have to be used they should not be used for other network traffic that could flood the switches and cause heartbeat delays.
 - If you cannot use a dedicated switch, the use of a dedicated VLAN can help limit the broadcast domain to protect the heartbeat traffic and the bandwidth it creates.
- For clusters of three or four FortiGate units, use switches to connect heartbeat interfaces. The corresponding heartbeat interface of each FortiGate unit in the cluster must be connected to the same switch. For improved redundancy use a different switch for each heartbeat interface. In that way if the switch connecting one of the heartbeat interfaces fails or is unplugged, heartbeat traffic can continue on the other heartbeat interfaces and switch.
- Isolate heartbeat interfaces from user networks. Heartbeat packets contain sensitive cluster configuration information and can consume a considerable amount of network bandwidth. If the cluster consists of two FortiGate units, connect the heartbeat interfaces directly using a crossover cable or a regular Ethernet cable. For clusters with more than two units, connect heartbeat interfaces to a separate switch that is not connected to any network.
- If heartbeat traffic cannot be isolated from user networks, enable heartbeat message encryption and authentication to protect cluster information. See [Enabling or disabling HA heartbeat encryption and authentication on page 1529](#).
- Configure and connect redundant heartbeat interfaces so that if one heartbeat interface fails or becomes disconnected, HA heartbeat traffic can continue to be transmitted using the backup heartbeat interface. If heartbeat communication fails, all cluster members will think they are the primary unit resulting in multiple devices on the network with the same IP addresses and MAC addresses (condition referred to as *Split Brain*) and communication will be disrupted until heartbeat communication can be reestablished.
- Do not monitor dedicated heartbeat interfaces; monitor those interfaces whose failure should trigger a device failover.
- Where possible at least one heartbeat interface should not be connected to an NP4 or NP6 processor to avoid NP4 or NP6-related problems from affecting heartbeat traffic.
- Where possible, the heartbeat interfaces should not be connected to an NP4 or NP6 processor that is also processing network traffic.
- Where possible, each heartbeat interface should be connected to a different NP4 or NP6 processor.
- Any FortiGate interface can be used as a heartbeat interface including 10/100/1000Base-T, SFP, QSFP fiber and copper, and so on. If you set up two or more interfaces as heartbeat interfaces each interface can be a different type and speed.

Interface monitoring (port monitoring)

Fortinet suggests the following practices related to interface monitoring (also called port monitoring):

- Wait until a cluster is up and running and all interfaces are connected before enabling interface monitoring. A monitored interface can easily become disconnected during initial setup and cause failovers to occur before the cluster is fully configured and tested.
- Monitor interfaces connected to networks that process high priority traffic so that the cluster maintains connections to these networks if a failure occurs.
- Avoid configuring interface monitoring for all interfaces.
- Supplement interface monitoring with remote link failover. Configure remote link failover to maintain packet flow if a link not directly connected to a cluster unit (for example, between a switch connected to a cluster interface and the network) fails. See [Remote link failover on page 1556](#).

WAN Optimization

WAN Optimization features require significant memory resources and generate a high amount of I/O on disk. Before enabling WAN Optimization, ensure that the memory usage is not too high. If possible, avoid other disk-intensive features such as heavy traffic logging on the same disk as the one configured for WAN Optimization needs.

In general, it is preferable to enable the Transparent Mode checkbox and ensure that routing between the two endpoints is acceptable. Some protocols may not work well without enabling Transparent Mode.

Other best practices for utilizing the WAN Optimization feature follow.

Sharing the WAN Opt. tunnel for traffic of the same nature

WAN optimization tunnel sharing is recommended for similar types of WAN optimization traffic (such as CIFS traffic from different servers). However, tunnel sharing for different types of traffic is not recommended. For example, aggressive and non-aggressive protocols should not share the same tunnel.

Ordering WAN Opt. rules appropriately

- Precise, port specific WAN Optimization rules should be at the top of the list.
- Generic rules, such as overall TCP, should be at the bottom of the list.

Avoiding mixing protocols in a WAN Opt. tunnel

Different protocols may be more or less talkative or interactive. Mixing protocols in a tunnel may result in a delay for some of them. It is recommended to define protocol specific wan-optimization rules and restrict the ports to the necessary ones only for performance reasons.

Setting correct configuration options for CIFS WAN Opt.

Ensure that the WAN Optimization rules cover TCP ports 139 and 445 (on the same or two different rules). Also ensure that Transparent Mode is selected.

Setting correct configuration options for MAPI WAN Opt.

For MAPI WAN Optimization, only specify a rule with TCP port 135 (unless the MAPI control port is configured differently). Derived data sessions using other random ports will be handled by the CIFS wan-optimization daemon even with only the control port configured.

Testing WAN Opt. in a lab

- Ensure that WAN emulators are used to simulate the WAN. If no WAN emulator is used, it is expected to have better results without WAN Optimization than with WAN Optimization.
- To test the difference between cold transfers (first-time transfers) and warm transfers, it is recommended to generate a random file of the cold transfer to ensure that the test is the first time that the file has been seen.

Regarding byte compression and type of file

Enabling byte compression on file transfers already compressed (.jpeg files, compressed archive, etc.) won't provide any performance increase and could be seen as a misuse of CPU resources.

Regarding network address translation (NAT)

Selecting the NAT feature in a security policy does not have any influence on WAN Optimization traffic.

High Availability

There is no benefit to using active-active mode, so for pure WAN Optimization needs, use active-passive mode. Refer to the [FGCP high availability](#) section for other best practices related to HA.

Authentication with specific peers

Configure WAN optimization authentication with specific peers. Accepting any peer is not recommended as this can be less secure.

Virtual Domains (VDOMs)

VDOMs can provide separate firewall policies and, in NAT/Route mode, completely separate configurations for routing and VPN services for each connected network or organization. This section provides a list of best practices for configuring VDOMs.

Per-VDOM resource settings

While Global resources apply to resources shared by the whole FortiGate unit, per-VDOM resources are specific to only one Virtual Domain.

By default all the per-VDOM resource settings are set to no limits. This means that any single VDOM can use up all the resources of the entire FortiGate unit if it needs to do so. This would starve the other VDOMs for resources to the point where they would be unable to function. For this reason, it is recommended that you set some maximums on resources that are most vital to your customers.

Virtual domains in NAT/Route mode

Once you have enabled virtual domains and created one or more VDOMs, you need to configure them. It is recommended that you perform the following tasks in the order given (while you may not require all for your network topology):

1. Change the management virtual domain.
2. Configure FortiGate interfaces in a NAT/Route VDOM.
3. Configure VDOM routing.
4. Configure security policies for NAT/Route VDOMs.
5. Configure UTM profiles for NAT/Route VDOMs.
6. Test the configuration.

Virtual clustering

If you decide to disable override for clustering, as a result of persistent renegotiating, you should disable it for both cluster units.

Explicit proxy

- For explicit proxies, when configuring limits on the number of concurrent users, you need to allow for the number of users based on their authentication method. Otherwise you may run out of user resources prematurely.
- Each session-based authenticated user is counted as a single user using their authentication membership (RADIUS, LDAP, FSSO, local database etc.) to match users in other sessions. So one authenticated user in multiple sessions is still one user.
- For all other situations, the source IP address is used to determine a user. All sessions from a single source address are assumed to be from the same user.
- Set the explicit web proxy and explicit FTP proxy Default Firewall Policy Action to Deny. This means that a firewall policy is required to use these explicit proxies, allowing you to control access and impose security features.
- Do not enable the explicit web or FTP proxy on an interface connected to the Internet. This is a security risk because anyone on the Internet who finds the proxy could use it to hide their source address. If you must enable the proxy on such an interface make sure authentication is required to use the proxy.

Wireless

The following section contains a list of best practices for wireless network configurations with regard to encryption and authentication, geographic location, network planning, power usage, client load balancing, local bridging, SSIDs, and the use of static IPs.

Encryption and authentication

It is best practice to always enable the strongest user authentication and encryption method that your client supports. Fortinet recommends the following security, in order of strongest to weakest:

- WPA2 - Enterprise 802.1x/EAP - Personal pre-shared key (8-63 characters)
- WPA - Enterprise 802.1x/EAP - Personal pre-shared key (8-63 characters)
- WEP128 - 26 Hexadecimal digit key
- WEP64 - 10 Hexadecimal digit key
- None - Open system

Geographic location

Ensure that the FortiGate wireless controller is configured for your geographic location. This ensures that the available radio channels and radio power are in compliance with the regulations in your region.

The maximum allowed transmitter power and permitted radio channels for Wi-Fi networks depend on the region in which the network is located. By default, the WiFi controller is configured for the United States. If you are located in any other region, you need to set your location before you begin configuring wireless networks.

The location setting can only be changed from CLI. To change the country to France, for example, enter the following:

```
config wireless-controller setting
  set country FR
end
```

To see the list of country codes, enter a question mark ('?') in place of the country code.

Using an incorrect geographic location is a common error that can lead to unpredictable results on the client side.

Network planning

It is recommended that you perform a proper site survey prior positioning the wireless access point. In order to evaluate the coverage area environment, the following criteria must be taken into account:

- Size of coverage area
- Bandwidth required
- Client wireless capabilities

After completing a RF site survey, you'll have a good idea of the number and location of access points needed to provide users with adequate coverage and performance.

However, prior to installing the access points, be sure to determine the RF channel(s) you plan to use. This will ensure that users can roam throughout the facility with substantial performance.

To avoid co-channel interference, adjacent Wi-Fi APs must be configured to use non-overlapping channels. Otherwise, you'll find poor performance will degrade because of interference between access points.

It is recommended to statically configure the non-overlapping channels on every access point, using one Custom AP profile per AP (or group of APs). If static configuration cannot be used, the FortiOS Wi-Fi Controller includes the Automatic Radio Resource Provisioning (ARRP) feature.

Lowering the power level to reduce RF interference

Relevant Product(s): FortiAP

Reducing power reduces unwanted coverage and potential interference to other WLANs. Areas of unwanted coverage are a potential security risk. If possible, reduce the transmitter power of your wireless access point so that the signal is not available beyond the areas where it is needed. Auto Tx Power Control can be enabled to automatically adjust the transmit power.

In cases where customers complain about slow wireless traffic through a FortiAP, it might be necessary to try to reduce the possibility of RF interference. It is best practice not to locate FortiAPs near steel beams or other interfering materials. You can try using a wireless sniffer tool to collect the wireless packets and then analyze the extent of air interference.

A common mistake is spacing FortiAPs based upon the 5Ghz radio frequency. The 2.4Ghz signal travels further.

You have two options when confronted with slow wireless traffic through a FortiAP:

Option #1: Reducing transmit power

Perform a speed test and record the results. Set one of the radios on a FortiAP to be in dedicated monitoring mode. Then observe how many APs are detected. If the number of APs is too high (i.e., greater than 20), try reducing the transmit power in the WTP profile for the FortiAPs until the number of dedicated APs has dropped significantly.

Repeat the speed test.

Option #2: Ensuring that VAPs are distributed over the available channels

No built-in tools are available to measure RF interference directly. However, FortiOS 5.0 does allow for automatic power adjustment, which should minimize the occurrence of RF interference.

Wireless client load balancing

Wireless load balancing allows your wireless network to more efficiently distribute wireless traffic among wireless access points and available frequency bands. FortiGate wireless controllers support the following types of client load balancing:

- **Access Point Hand-off** - The wireless controller signals a client to switch to another access point.
- **Frequency Hand-off** - The wireless controller monitors the usage of 2.4GHz and 5GHz bands, and signals clients to switch to the lesser-used frequency.

Local bridging

Whenever possible, use local bridging to offload the CAPWAP tunnel. Note that in this case, Wi-Fi client devices obtain IP addresses from the same DHCP server as wired devices on the LAN. The vlan ID can only be configured

from the CLI:

```
config wireless-controller vap
  edit "vaplocalbridge"
    set vdom "root"
    set ssid "testvaplocalbridge"
    set local-bridging enable
    set vlanid 40 ---> only available in CLI
  next
end
```

Advertising SSIDs

- It is highly recommended to advertise the SSID. It makes it easier for customers and wireless clients. Also, if you 'hide' the SSID (known as 'network cloaking'), then clients will always look for it when they're outside the coverage area, which searches for known SSIDs, in effect leaking the SSID anyway. Refer to [RFC 3370](#). Furthermore, many of the latest Broadcom drivers do not support hidden SSID for WPA2.
- For security reason, you might want to prevent direct communication between your wireless clients. In this case, enable Block Intra-SSID Traffic (in the SSID configuration).
- In a network with multiple wireless controllers, you need to change the mesh SSID so that each mesh root has a unique SSID. Other controllers using the same mesh root SSID might be detected as fake or rogue APs. Go to **WiFi & Switch Controller > SSID** to change the SSID. Fortinet also recommends that you create a new preshared key instead of using the default.

Using static IPs in a CAPWAP configuration

In a large FortiAP deployment with more than 20 FortiAPs connecting to a Fortigate Wireless Controller (AC), it is recommended to use static IPs on the access points instead of DHCP, setting the AC IP statically and the AC discovery type to static (Type 1), instead of learning it through broadcast, multicast, or DHCP.

This makes management of the APs easier since you know the exact IP of each access point. Troubleshooting also becomes easier as the debug of the AC controller won't continuously attempt the different discovery methods in sequence (broadcast > multicast > static).

Logging and reporting

The default log device settings must be modified so that system performance is not compromised. The FortiGate unit, by default, has all logging of FortiGate features enabled, except for traffic logging. The default logging location will be either the FortiGate unit's system memory or hard disk, depending on the model. Units with a flash disk are not recommended for disk logging.

Log management

When the FortiGate unit records FortiGate activity, valuable information is collected that provides insight into how to better protect network traffic against attacks, including misuse and abuse. There is a lot to consider before enabling logging on a FortiGate unit, such as what FortiGate activities to enable and which log device is best suited for your network's logging needs. A plan can help you in deciding the FortiGate activities to log, a log device, as well as a backup solution in the event the log device fails.

This plan should provide you with an outline, similar to the following:

- What FortiGate activities you want and/or need logged (for example, security features).
- The logging device best suited for your network structure.
- If you want or require archiving of log files.
- Ensuring logs are not lost in the event a failure occurs.

After the plan is implemented, you need to manage the logs and be prepared to expand on your log setup when the current logging requirements are outgrown. Good log management practices help you with these tasks.

Log management practices help you to improve and manage logging requirements. Logging is an ever-expanding tool that can seem to be a daunting task to manage. The following management practices will help you when issues arise, or your logging setup needs to be expanded.

- Revisit your plan on a yearly basis to verify that your logging needs are being met by your current log setup. For example, your company or organization may require archival logging, but not at the beginning of your network's lifespan. Archival logs are stored on a FortiGate unit's local hard drive, a FortiAnalyzer unit, or a FortiCloud server, in increasing order of size.
- Configure an alert message that will notify you of activities that are important to be aware about. For example: if a branch office does not have a FortiGate administrator, you will need to know at all times that the IPsec VPN tunnel is still up and running. An alert email notification message can be configured to send only if IPsec tunnel errors occur.
- If your organization or company uses peer-to-peer programs such as Skype or other instant messaging software, use the IM usage dashboard widget or the Executive Summary's report widget (Top 10 Application Bandwidth Usage Per Hour Summary) to help you monitor the usage of these types of instant messaging software. These widgets can help you in determining how these applications are being used, including if there is any misuse and abuse. Their information is taken from application log messages; however, application log messages should be viewed as well since they contain the most detailed information.
- Ensure that your backup solution is up-to-date. If you have recently expanded your log setup, you should also review your backup solution. The backup solution provides a way to ensure that all logs are not lost in the event that the log device fails or issues arise with the log device itself.

- When downloading log messages and viewing them on a computer, the log file will be downloaded like any other file. Log file names contain their log type and date in the name, so it is recommended to create a folder in which to archive your log messages, as they can be sorted easily.

System memory and hard disks

If the FortiGate unit has a hard disk, it is enabled by default to store logs. This also means that you do not have to enable this and configure the settings for logging to the hard disk, but modify these settings so that it is configured for your network logging requirements.

If the FortiGate unit has only flash memory, disk logging is disabled by default, as it is not recommended. Constant rewrites to flash drives can reduce the lifetime and efficiency of the memory. It must be enabled in the CLI under config log disk setting.

For some low-end models, disk logging is unavailable. Check a product's Feature Matrix for more information. In either case, Fortinet recommends using either a FortiAnalyzer unit or the FortiCloud service.

Chapter 5 - FortiOS Carrier

This FortiOS Handbook chapter contains the following sections:

- [What's new in FortiOS Carrier 6.0](#)
 - Features and improvements that have been added in version FortiOS 5.6
- [Overview of FortiOS Carrier features](#)
 - The high altitude overview of FortiOS Carrier
- ["MMS Concepts" on page 382](#)
 - Basic background on MMS and its concepts so that there is some context to the settings used in the MMS configuration
- ["MMS Configuration" on page 397](#)
 - Listing of the MMS configuration options and settings
 - Procedures and insights in the configuration of the MMS components of FortiOS Carrier
- [GTP basic concepts](#)
 - Basic background on GTP and its concepts so that there is some context to the settings used in the GTP configuration
- ["GTP Configuration" on page 462](#)
 - Listing of the GTP configuration options and settings
 - Procedures and insights in the configuration of the GTP components of FortiOS Carrier
- ["SCTP Concepts" on page 506](#)
 - Background on SCTP, and how it makes the FortiOS Carrier different than the FortiGate
- ["Troubleshooting" on page 510](#)
 - Some basic procedures for troubleshooting FortiOS Carrier

What's new in FortiOS Carrier 6.0

The following list contains new firewall features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- [FGSP support over GTP on page 1](#)
- [GTP Stats via SNMP on page 488](#)
- [GTP Shared Tunnel Limit on page 486](#)

The GTP Shared Tunnel limit is the total number of GTP tunnels created by multiple GTP profiles. A Global shared tunnel limit gives the flexibility of limiting the number of GTP tunnels flowing through different profiles. A shared global limit is defined and then referenced in the profiles. Before FortiOS version 6, the GTP tunnel limit could be set on a per-profile basis. The GTP tunnel limits can now be set per VDOM.

Per profile tunnel limiting is still possible but restrictive limit between the global limit and the profile limit will be enforced.

Example:

- Global shared tunnel limit defined as 12
- GTP Profile A - per profile tunnel limit defined as 8
- GTP Profile B - per profile tunnel limit defined as 14

You can have eight tunnels active in profile A, but the ninth will be dropped due to the profile limit of 8.

If Profile A still has 8 active tunnels, you can have four tunnels active in profile B and the fifth will be dropped even though the profile allows 14, because the global share limit is 12.

```
next
edit "gtp2"
  set global-tunnel-limit "gtp-tl-1"
end
```

Overview of FortiOS Carrier features

FortiOS Carrier specific features include Multimedia messaging service (MMS) protection, and GPRS Tunneling Protocol (GTP) protection.

All FortiGate units, carrier-enabled or not, are capable of handling Stream Control Transmission Protocol (SCTP) traffic, which is a protocol designed for and primarily used in Carrier networks.

This section includes:

Overview

FortiOS Carrier provides all the features found on FortiGate units plus added features specific to carrier networks: MMS and GTP.

MMS

MMS is a standard for sending messages that include multimedia content between mobile phones. MMS is also popular as a method of delivering news and entertainment content including videos, pictures, and text. Carrier networks include four different MMS types of messages — MM1, MM3, MM4, and MM7.

GTP

The GPRS Tunneling Protocol (GTP) runs on GPRS carrier networks. GPRS is a GSM packet radio standard. It provides more efficient usage of the radio interface so that mobile devices can share the same radio channel. FortiOS supports GTPv1 and GTPv2.

GPRS provides direct connections to the Internet (TCP/IP) and X.25 networks for point-to-point services (connection-less/connection oriented) and point-to-multipoint services (broadcast).

GPRS currently supports data rates from 9.6 kbps to more than 100 kbps, and it is best suited for burst forms of traffic. GPRS involves both radio and wired components. The mobile phone sends the message to a base station unit (radio based) that converts the message from radio to wired, and sends the message to the carrier network and eventually the Internet (wired carrier network). See [GTP](#).


MMS Concepts

MMS background

MMS is a common method for mobile users to send and receive multimedia content. A Carrier network supports MMS across its network. This makes up the MMS Service Provider Network (MSPN).

Messages can be sent or received between the MMSC and a number of other services including the Internet, content providers, or other carriers. Each of these different service connections uses different MMS formats including MM1 and MM7 messages (essentially HTTP format), and MM3 and MM4 messages (SMTP formatted). These different formats reflect the different purposes and content for each type of MMS message.

MMS content interfaces



MMS content interfaces

MMS messages are sent from devices and servers to other devices and servers using MMS content interfaces

There are eight interfaces defined for the MMS standard, referred to as MM1 through MM8. The most important of these interfaces for the transfer of data is the MM1 interface, as this defines how mobile users communicate from the mobile network to the Multimedia Message Service Center (MMSC). MMS content to be monitored and controlled comes from these mobile users and is going to the provider network.

Other MMS content interfaces that connect a service provider network to other external sources can pose threats as well. MM3 handles communication between the Internet and the MMSC and is a possible source of viruses and other content problems from the Internet. MM4 handles communication between different content provider MMSCs. Filtering MM4 content protects the service provider network from content sent from foreign service providers and their subscribers. Finally MM7 is used for communication between content providers and the MMSC. Filtering MM3 content can also keep harmful content off of the service provider network.

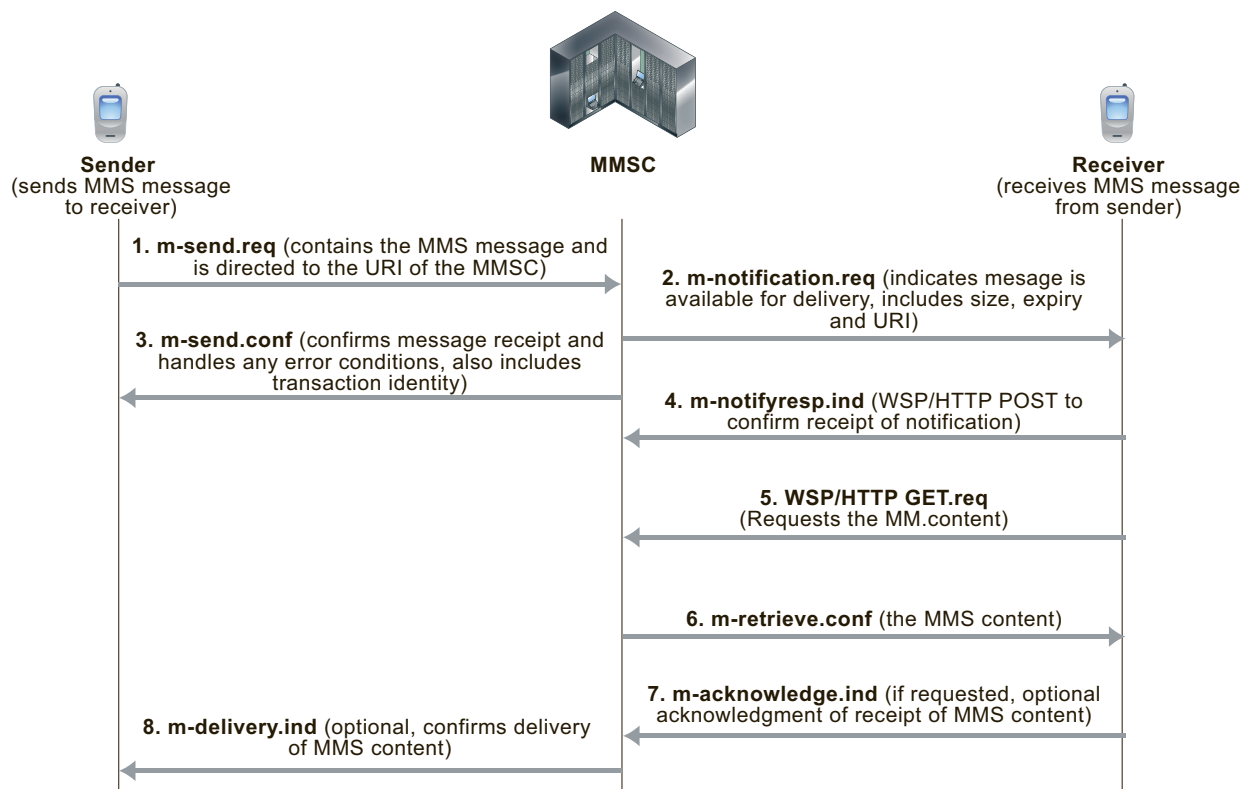
MMS content interfaces

Type	Transaction	Similar to
MM 1	Handset to MMSC	HTTP
MM 3	Between MMSC and Internet	SMTP
MM 4	Between Operator MMSCs	SMTP
MM 7	Content Providers to MMSC	HTTP and SOAP

How MMS content interfaces are applied

As shown below, the sender's mobile device encodes the MMS content in a form similar to MIME email message (MMS MIME content formats are defined by the MMS Message Encapsulation specification). The encoded message is then forwarded to the service provider's MMSC. Communication between the sending device and the MMSC uses the MM1 content interface. The MM1 content interface establishes a connection and sends an MM1 send request (`m-send.req`) message that contains the MMS message. The MMSC processes this request and sends back an MM1 send confirmation (`m-send.conf`) HTTP response indicating the status of the message — accepted or an error occurred, for example.

MM1 transactions between senders and receivers and the MMSC

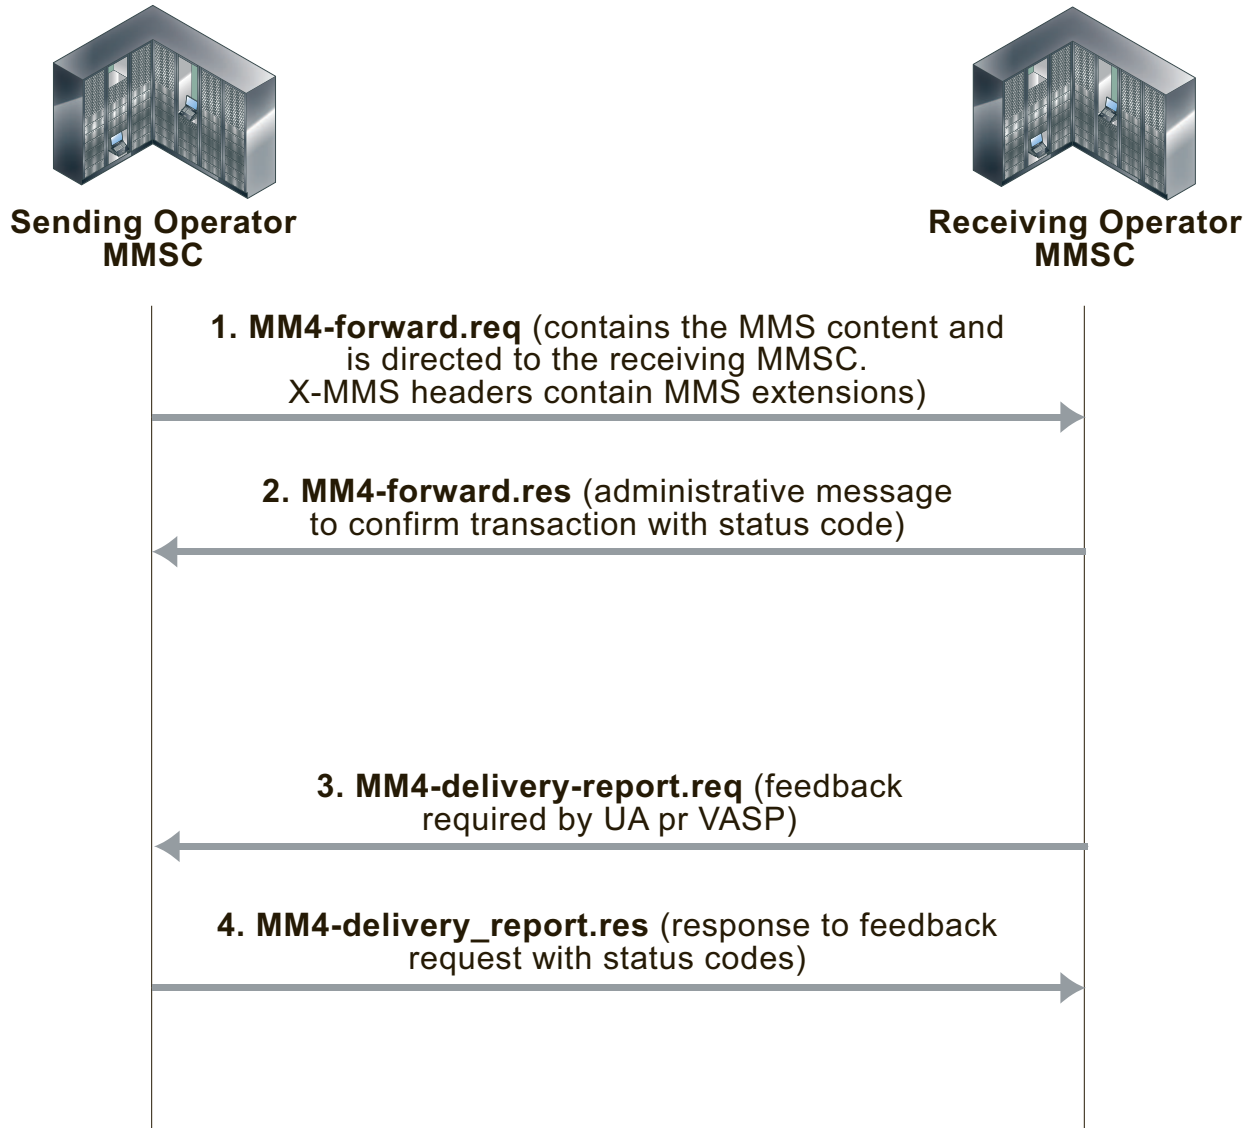


If the recipient is on another carrier, the MMSC forwards the message to the recipient's carrier. This forwarding uses the MM4 content interface for forwarding content between operator MMSCs (see the figure below).

Before the MMSC can forward the message to the final recipient, it must first determine if the receiver's handset can receive MMS messages using the MM1 content interface. If the recipient can use the MM1 content interface, the content is extracted and sent to a temporary storage server with an HTTP front-end.

To retrieve the message, the receiver's handset establishes a connection with the MMSC. An HTTP get request is then sent from the recipient to the MMSC. This message contains the URL where the content of the message is stored. The MMSC responds with a retrieve confirmation (`m-retrieve.conf`) HTTP response that contains the message.

MM4 messages sent between operator MMSCs



This causes the receiver's handset to retrieve the content from the embedded URL. Several messages are exchanged to indicate status of the delivery attempt. Before delivering content, some MMSCs also include a content adaptation service that attempts to modify the multimedia content into a format suitable for the recipient's handset.

If the receiver's handset is not MM1 capable, the message can be delivered to a web based service and the receiver can view the content from a normal Internet browser. The URL for the content can be sent to the receiver in an SMS text message. Using this method, non-MM1 capable recipients can still receive MMS content.

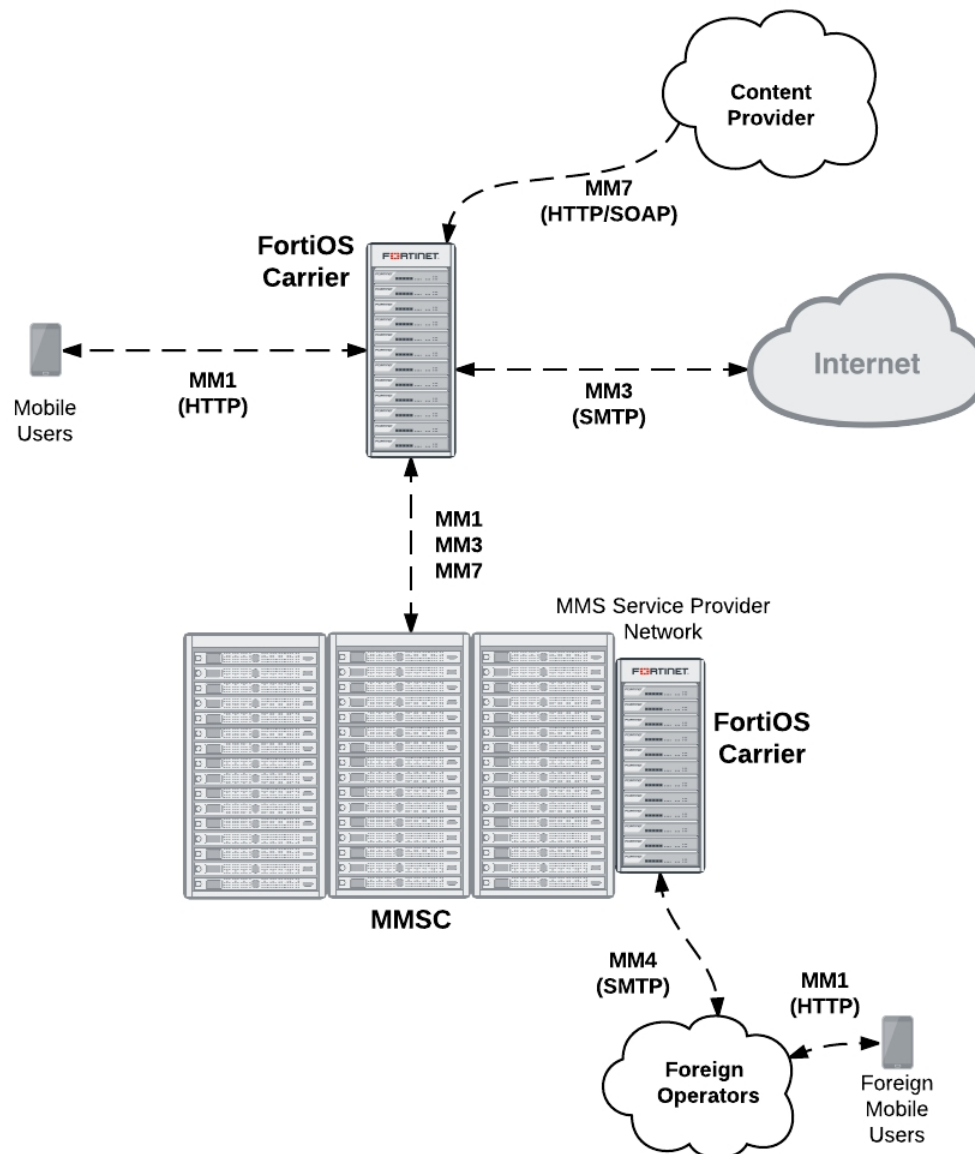
The method for determining whether a handset is MMS capable is not specified by the standards. A database is usually maintained by the operator, and in it each mobile phone number is marked as being associated with a legacy handset or not. It can be a bit hit and miss since customers can change their handset at will and this database is not usually updated dynamically.

Email and web-based gateways from MMSC to the Internet use the MM3 content interface. On the receiving side, the content servers can typically receive service requests both from WAP and normal HTTP browsers, so delivery via the web is simple. For sending from external sources to handsets, most carriers allow MIME encoded message to be sent to the receiver's phone number with a special domain.

How FortiOS Carrier processes MMS messages

MMS messages can be vectors for propagating undesirable content such as spam and viruses. FortiOS Carrier can scan MMS messages sent using the MM1, MM3, MM4, and MM7 content interfaces. You can configure FortiOS Carrier to scan MMS messages for spam and viruses by configuring and adding MMS protection profiles and adding the MMS protection profiles to security policies. You can also use MMS protection profiles to apply content blocking, carrier endpoint filtering, MMS address translation, sending MMS notifications, DLP archiving of MMS messages, and logging of MMS message activity.

FortiOS Carrier MMS processing



FortiOS Carrier can send MMS messages to senders informing those senders that their devices are infected. FortiOS Carrier can also send MMS notifications to administrators to inform them of suspicious activity on their networks.

For message floods and duplicate messages, FortiOS Carrier does not send notifications to message senders but does send notifications to administrators and sends messages to sender handsets to complete MM1 and MM4 sessions.

Where MMS messaging uses the TCP/IP set of protocols, SMS text messaging uses the Signaling System Number 7 (SS7) set of protocols, which is not supported by FortiOS.

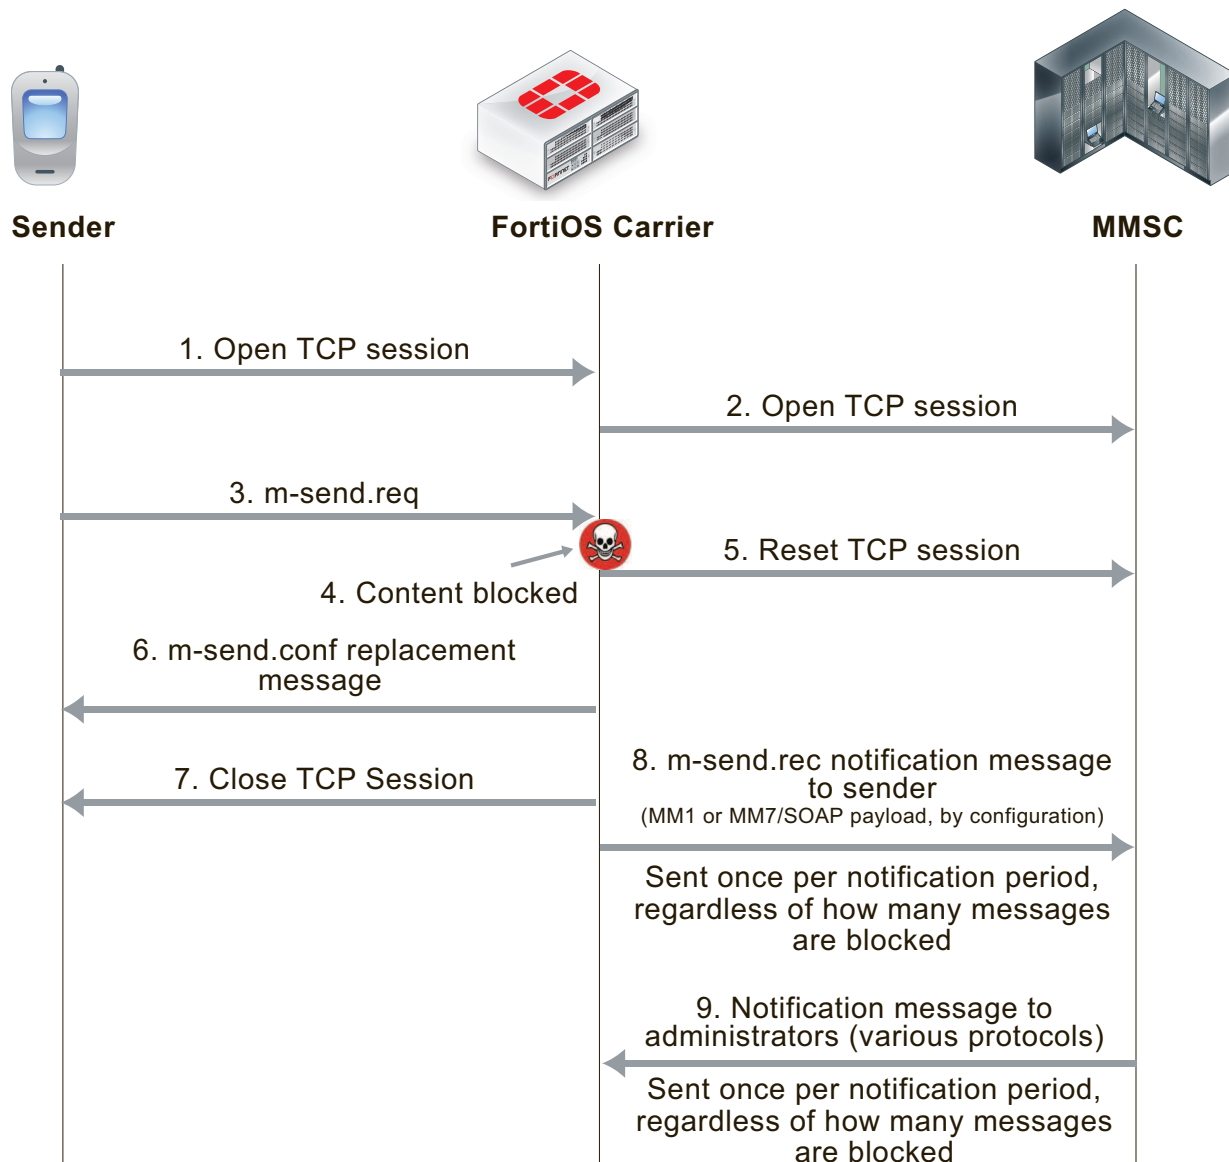
FortiOS Carrier and MMS content scanning

The following section applies to MMS content scanning, including virus scanning, file filtering, content spam filtering, carrier endpoint filtering, and MMS content checksum filtering.

MM1 Content Scanning

During MM1 content scanning a message is first transmitted from the sender, establishing a connection with the MMSC. FortiOS Carrier intercepts this connection and acts as the endpoint. FortiOS Carrier then establishes its own connection to the MMSC. Once connected, the client transmits its `m-send.req` HTTP post request to FortiOS Carrier which scans it according to the MMS protection profile settings. If the content is clean, the message is forwarded to the MMSC. The MMSC returns `m-send.conf` HTTP response through FortiOS Carrier to the sender.

If FortiOS Carrier blocks the message (for example because a virus was found, see the figure below), FortiOS Carrier resets the connection to the MMSC and sends `m-send.conf` HTTP response back to the sender. The response message can be customized using replacement messages. FortiOS Carrier then terminates the connection. Sending back an `m-send.conf` message prevents the sender from trying to send the message again.

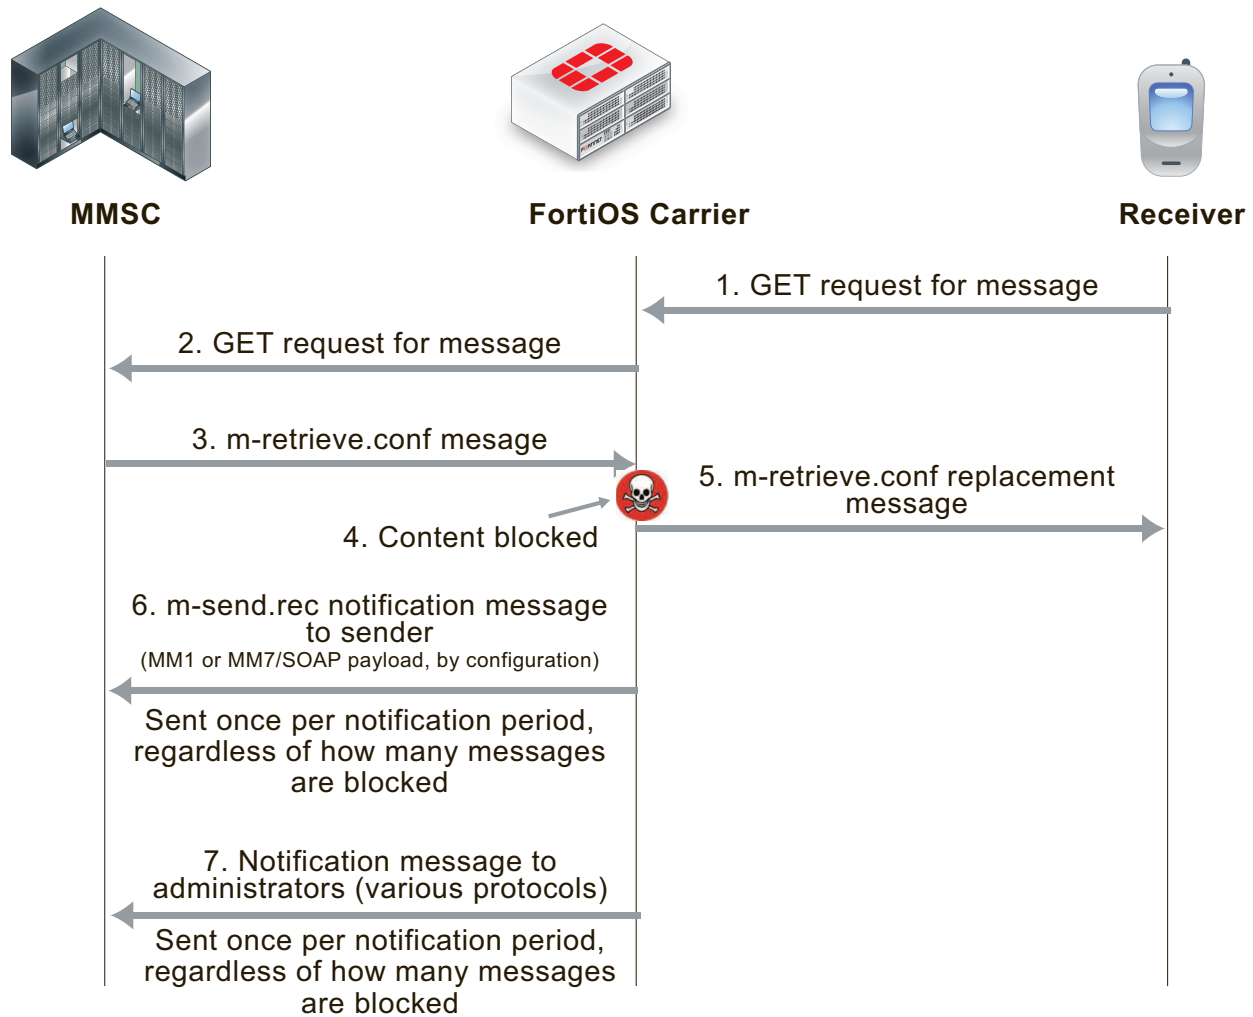
MM1 MMS scanning of message sent by sender (blocking m.send.req messages)

FortiOS Carrier also sends `m-send.rec` notifications messages to the MMSC that are then forwarded to the sender to notify them of blocked messages.

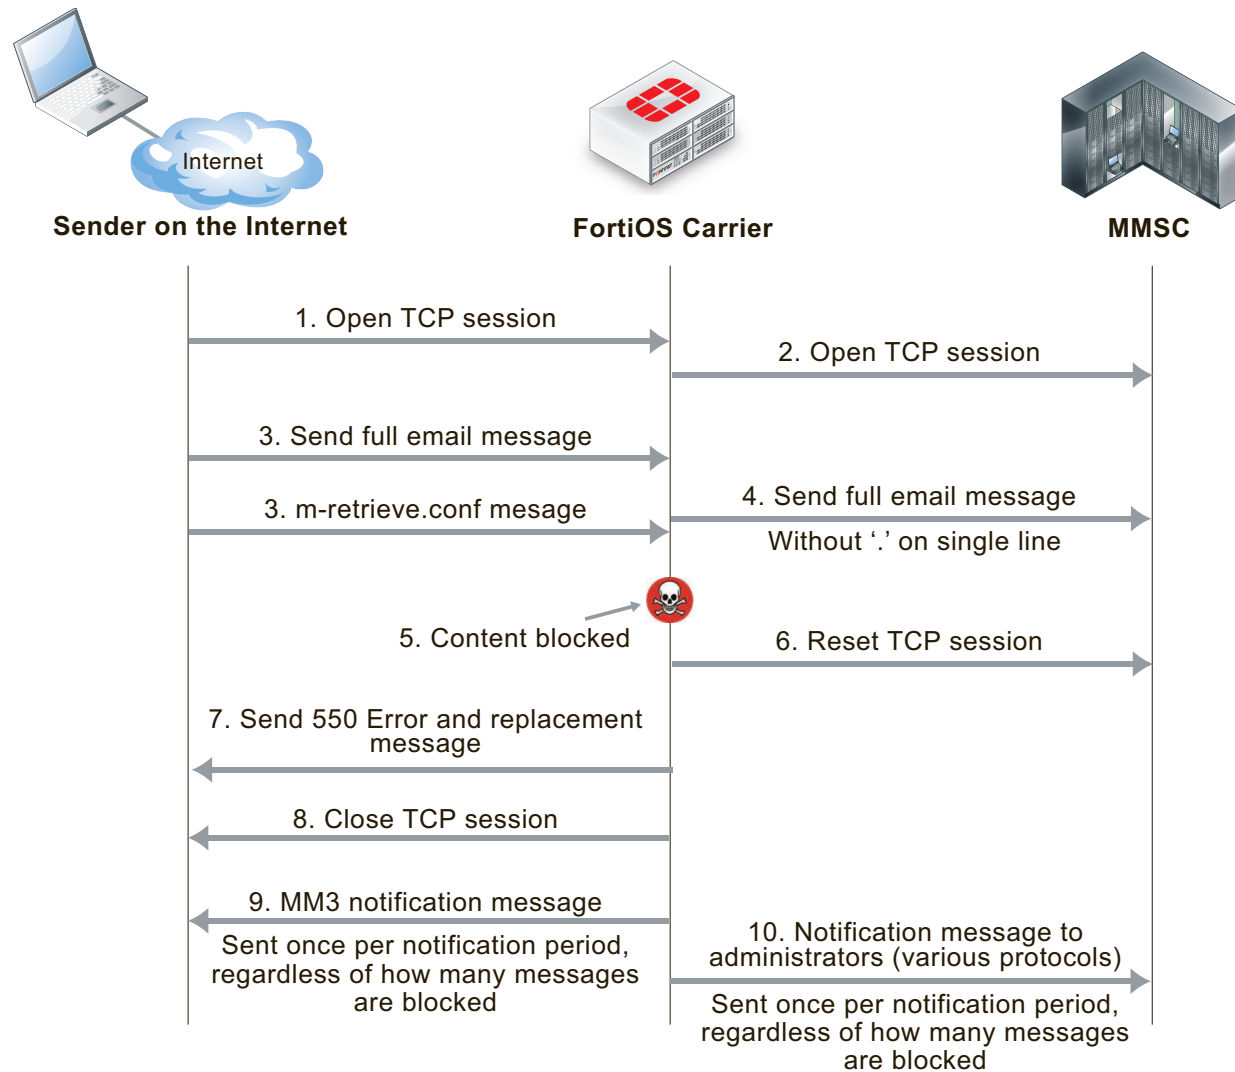
Filtering message retrieval

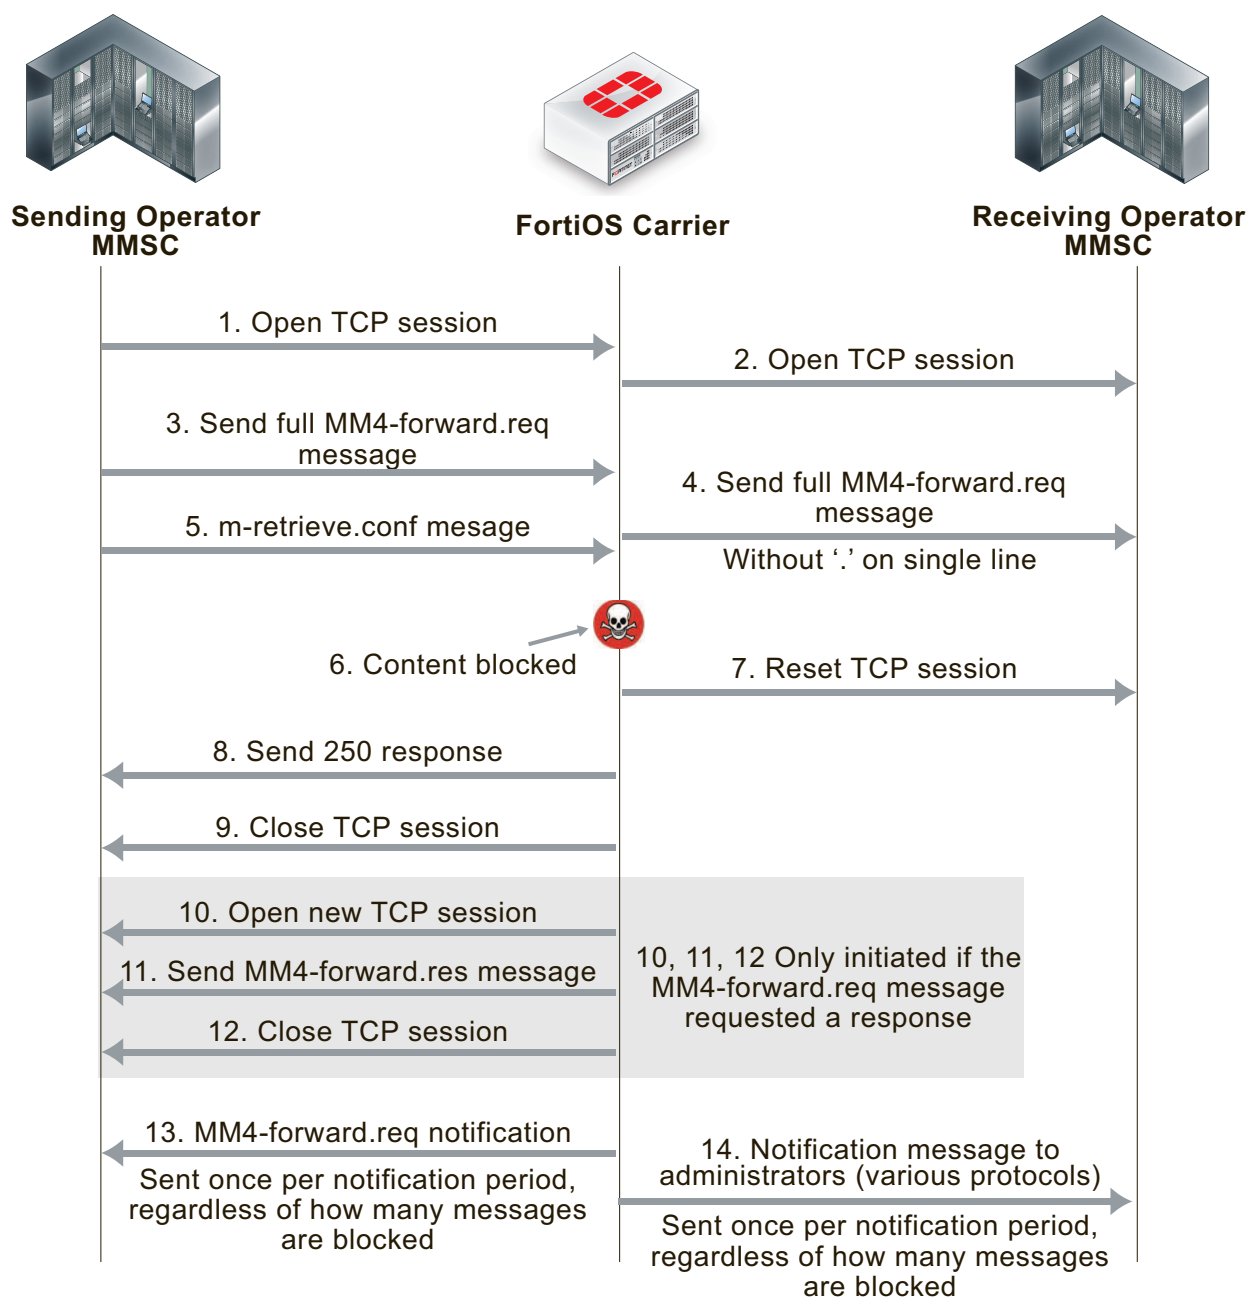
FortiOS Carrier intercepts the connection to the MMSC, and the `m-retrieve.conf` HTTP response from the MMSC is scanned according to the MMS content scanning settings. If the content is clean, the response is forwarded back to the client. If the content is blocked, FortiOS Carrier drops the connection to the MMSC. It then builds an `m-retrieve.conf` message from the associated replacement message and transmits this back to the client.

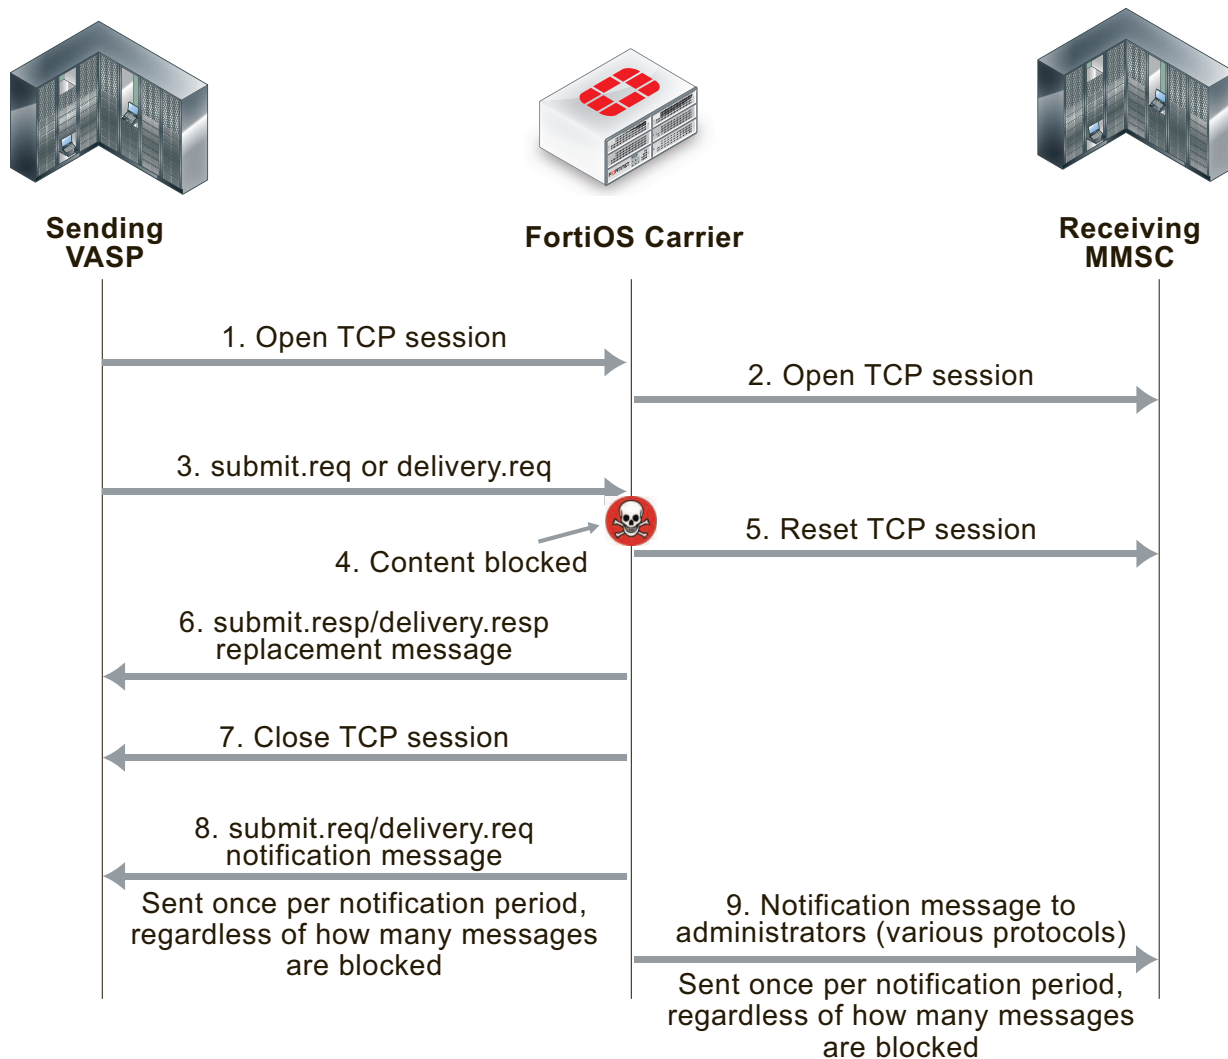
FortiOS Carrier also sends `m-send.rec` notifications messages to the MMSC that are then forwarded to the receiver to notify them of blocked messages.

MM1 MMS scanning of messages received by receiver (blocking m.retrieve.conf messages)

Filtering MM3 and MM4 messages works in an similar way to MM1 (see the figures below). FortiOS Carrier intercepts connections to the MMSC, and scans messages as configured. When messages are blocked, FortiOS Carrier closes sessions as required, sends confirmation messages to the sender, notifies administrators, and notifies senders and receivers of messages.

MM3 MMS scanning of messages sent from a sender on the Internet to an MMSC

MM4 MMS scanning of messages sent between operator MMSCs

MM7 MMS scanning of messages sent between a VASP and an MMSC**FortiOS Carrier and MMS duplicate messages and message floods**

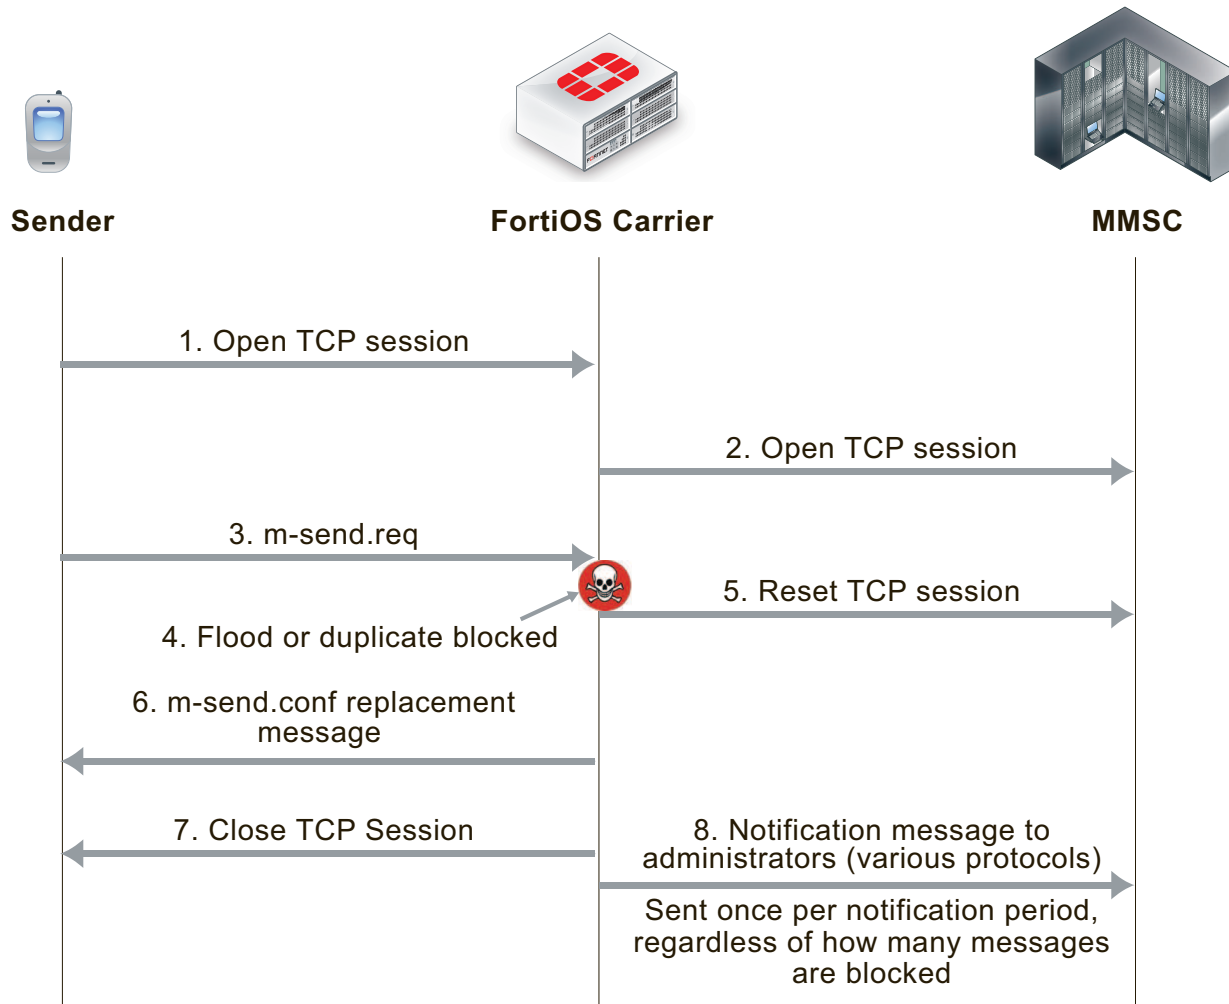
FortiOS Carrier detects duplicate messages and message floods for the MM1 and MM4 interfaces. How FortiOS Carrier detects and responds to duplicate messages and message floods is different from how FortiOS Carrier detects and responds to viruses and other MMS scanning protection measures.

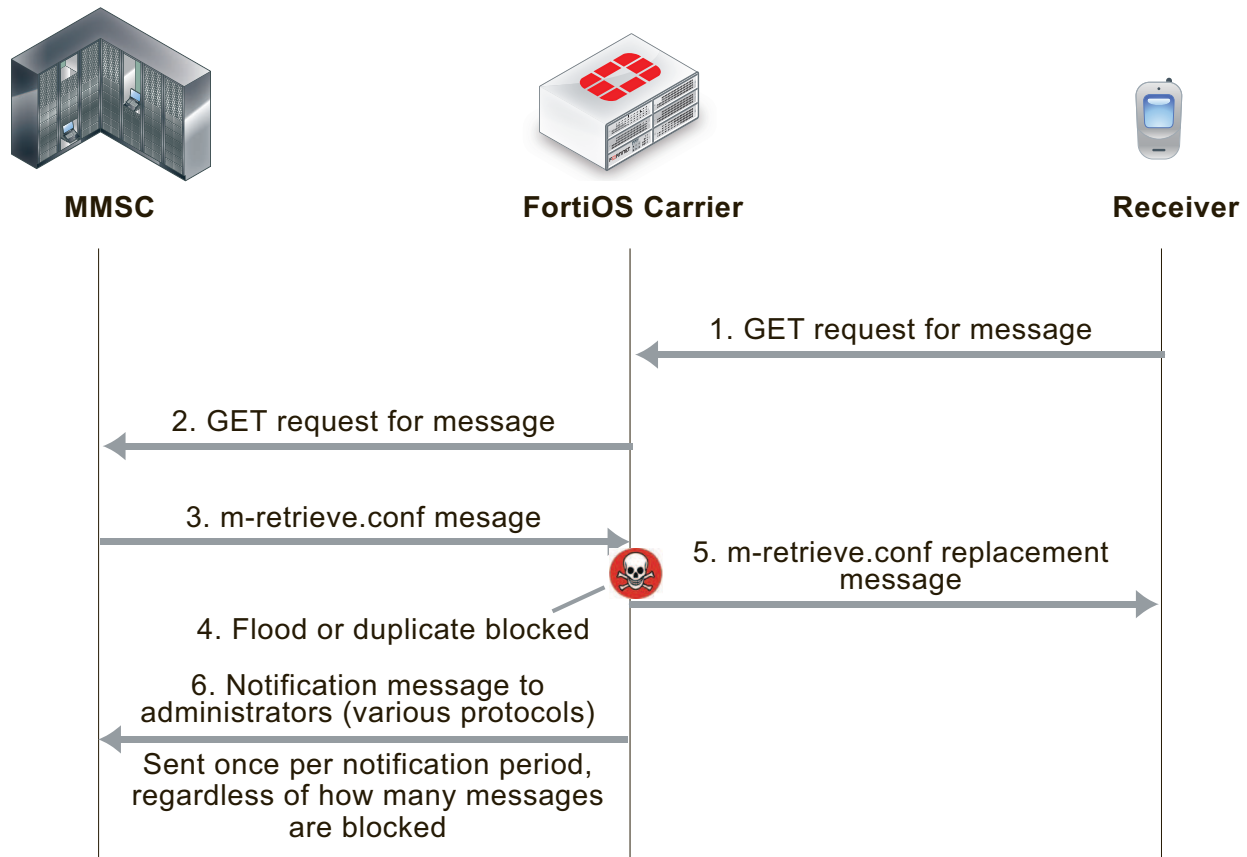
For message floods and duplicate messages, the sender does not receive notifications about floods or duplicate messages, as if the sender is an attacker they can gain useful information about flood and duplicate thresholds. Plus, duplicate messages and message floods are usually a result of a large amount of messaging activity and filtering of these messages is designed to reduce the amount of unwanted messaging traffic. Adding to the traffic by sending notifications to senders and receivers could result in an increase in message traffic.

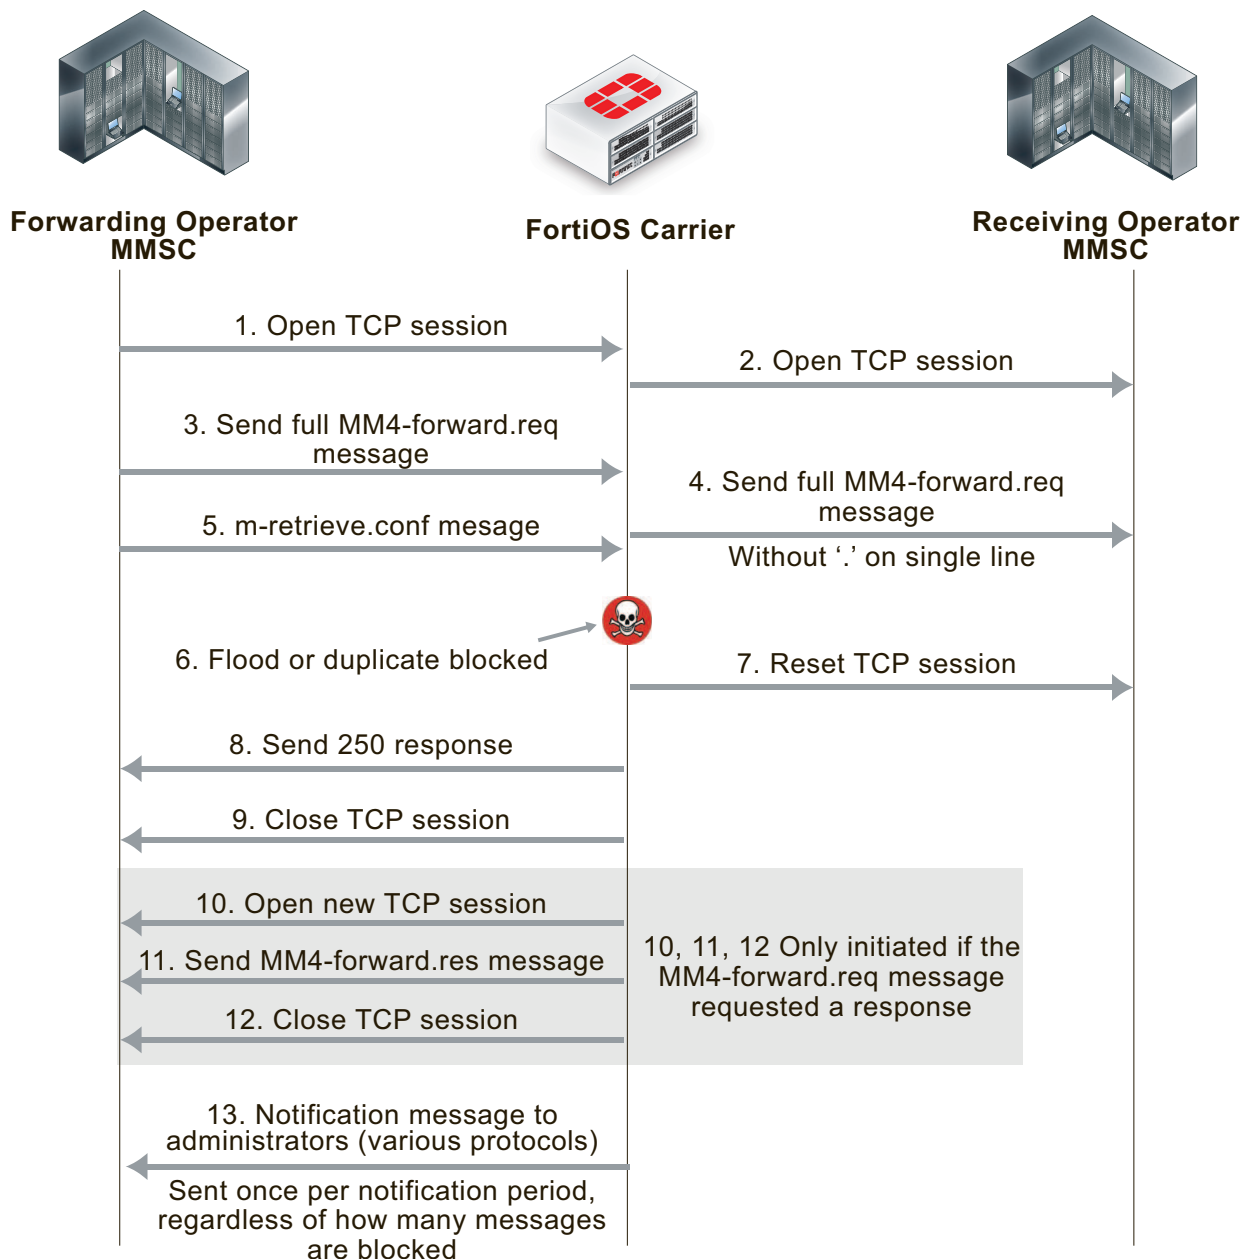
You can create up to three thresholds for detecting duplicate messages and message floods. For each threshold you can configure the FortiOS Carrier unit to respond by logging the activity, archiving or quarantining the messages, notifying administrators of the activity, and by blocking the messages. In many cases you may only want to configure blocking for higher activity thresholds, and to just monitor and send administrator notifications at lower activity thresholds.

When a block threshold is reached for MM1 messages, FortiOS Carrier sends `m-send.conf` or `m-retrieve.conf` messages to the originator of the activity. These messages are sent to end the MM1 sessions, otherwise the originator would continue to re-send the blocked message. When a block threshold is reached for MM4, FortiOS Carrier sends a `MM4-forward.res` message to close the MM4 session. An MM4 message is sent only if initiated by the originating `MM4-forward.req` message.

MM1 message flood and duplicate message blocking of sent messages



MM1 message flood and duplicate message blocking of received messages

MM4 message flood and duplicate message blocking**MMS protection profiles**

An MMS protection profile is a group of settings that you can apply to an MMS session matched by a security policy.

MMS protection profiles are easy to configure and can be used by more than one security policy. You can configure a single MMS protection profile for the different traffic types handled by a set of security policies that require identical protection levels and types. This eliminates the need to repeatedly configure those same MMS protection profile settings for each individual security policy.

For example, while traffic between trusted and untrusted networks might need strict protection, traffic between trusted internal addresses might need only moderate protection. You would configure two separate MMS protection profiles to provide the different levels of protection: one for traffic between trusted networks, and one for traffic between trusted and untrusted networks.

Once you have configured the MMS Protection Profile, you need to add it to a security policy to apply the profile to MMS traffic.

Bypassing MMS protection profile filtering based on carrier endpoints

You can use carrier endpoint filtering to exempt MMS sessions from MMS protection profile filtering. Carrier endpoint filtering matches carrier endpoints in MMS sessions with carrier endpoint patterns. If you add a carrier endpoint pattern to a filter list and set the action to exempt from all scanning, all messages from matching carrier endpoints bypass MMS protection profile filtering. See Bypassing message flood protection based on user's carrier endpoints.

Applying MMS protection profiles to MMS traffic

To apply an MMS protection profile you must first create the MMS protection profile and then add the MMS protection profile to a security policy by enabling the Carrier security profile. The MMS protection profile then applies itself to the traffic accepted by that security policy.

MMS protection profiles can contain settings relevant to many different services. Each security policy uses the subset of the MMS protection profile settings that apply to the sessions accepted by the security policy. In this way, you might define just one MMS protection profile that can be used by many security policies, each policy using a different or overlapping subset of the MMS protection profile.

To add an MMS protection profile to a security policy

1. Go to **Security Profiles > MMS Profile**.
2. Select **Create New** to add an MMS protection profile.
3. Configure as needed, and save.
4. Go to **Policy & Objects > IPv4 Policy**.
5. Select **Create New** to add a security policy, or select an existing policy and **Edit** to add the MMS profile.
6. Configure the security policy as required.
7. Enable **MMS Profile**, and select the MMS profile to add to the security policy.
8. Select **OK**.

MMS Configuration

MMS profiles

Since MMS profiles can be used by more than one security policy, you can configure one profile for the traffic types handled by a set of security policies requiring identical protection levels and types, rather than repeatedly configuring those same profile settings for each individual security policy.



If the security policy requires authentication, do not select the MMS profile in the security policy. This type of profile is specific to the authenticating user group. For details on configuring the profile associated with the user group, see User Groups in the Authentication guide.

For example, while traffic between trusted and untrusted networks might need strict protection, traffic between trusted internal addresses might need moderate protection. To provide the different levels of protection, you might configure two separate protection profiles: one for traffic between trusted networks, and one for traffic between trusted and untrusted networks.

Once you have configured the MMS profile, you can then apply the profile to MMS traffic by applying it to a security policy.

MMS profiles can contain settings relevant to many different services. Each security policy uses the subset of the MMS profile settings that apply to the sessions accepted by the security policy. In this way, you might define just one MMS profile that can be used by many security policies, each policy using a different or overlapping subset of the MMS profile.

The MMS Profile page contains options for each of the following:

- MMS scanning
- MMS Bulk Email Filtering Detection
- MMS Address Translation
- MMS Notifications
- DLP Archive
- Logging

MMS profile configuration settings

The following are MMS profile configuration settings in **Security Profiles > MMS Profile**.

MMS Profile page	
Lists each individual MMS profile that you created. On this page, you can edit, delete or create an MMS profile.	
Create New	Creates a new MMS profile. When you select Create New , you are automatically redirected to the New MMS Profile page.
Edit	Modifies settings within an MMS profile. When you select Edit , you are automatically redirected to the Edit MMS Profile.
Delete	Removes an MMS profile from the list on the MMS Profile page.
	To remove multiple MMS profiles from within the list, on the MMS Profile page, in each of the rows of the profiles you want removed, select the check box and then select Delete .
	To remove all MMS profiles from the list, on the MMS Profile page, select the check box in the check box column, and then select Delete .
Name	The name of the MMS profile.

Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (**Security Profiles > Antivirus**), 1 appears in **Ref.** .

To view the location of the referenced object, select the number in **Ref.**, and the Object Usage window appears displaying the various locations of the referenced object.

To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:

Ref.

View the list page for these objects – automatically redirects you to the list page where the object is referenced at.

Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page.

View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.

New MMS Profile page

Provides settings for configuring an MMS profile. This page also provides settings for configuring DLP archives and logging.

Profile Name	Enter a name for the profile.
Comments	Enter a description about the profile. This is optional.
MMS Scanning	Configure MMS Scanning options.
MMS Bulk Email Filtering Detection	Configure MMS Bulk Email options.
MMS Address Translation	Configure MMS Address Translation options.
MMS Notifications	Configure MMS Notification options.
DLP Archive	Configure DLP archive option.
Logging	Configure logging options.

MMS scanning options

You can configure MMS scanning protection profile options to apply virus scanning, file filtering, content filtering, carrier endpoint blocking, and other scanning to MMS messages transmitted using the MM1, MM3, MM4 and

MM7 protocols.

The following are the MMS Scanning options that are available within an MMS profile. You can create an MMS profile in **Security Profiles > MMS Profile** or edit an existing one. You must expand MMS Scanning to access the following options.

MMS Scanning section of the New MMS Profile page	
Monitor Only	<p>Select to cause the unit to record log messages when MMS scanning options find a virus, match a file name, or match content using any of the other MMS scanning options. Select this option to be able to report on viruses and other problems in MMS traffic without affecting users.</p> <p>Tip: Select Remove Blocked if you want the unit to actually remove content intercepted by MMS scanning options.</p>
Virus Scan	<p>Select to scan attachments in MMS traffic for viruses.</p> <p>Since MM1 and MM7 use HTTP, the oversize limits for HTTP and the HTTP antivirus port configuration also applies to MM1 and MM7 scanning.</p> <p>MM3 and MM4 use SMTP and the oversize limits for SMTP and the SMTP antivirus port configuration also applies to MM3 and MM4 scanning.</p>
Scan MM1 message retrieval	<p>Select to scan message retrievals that use MM1. If you enable Virus Scan for all MMS interfaces, messages are also scanned while being sent. In this case, you can disable MM1 message retrieval scanning to improve performance.</p>
Remove Blocked	<p>Select to remove blocked content from each protocol and replace it with the replacement message.</p> <p>Select Constant if the unit is to preserve the length of the message when removing blocked content, as may occur when billing is affected by the length of the message.</p> <p>Tip: If you only want to monitor blocked content, select Monitor Only.</p>
Content Filter	<p>Select to filter messages based on matching the content of the message with the words or patterns in the selected web content filter list.</p> <p>For information about adding a web content filter list, see the FortiGate CLI Reference.</p>
Carrier Endpoint Block	<p>Select to add Carrier Endpoint Filtering in this MMS profile. Select the carrier endpoint filter list to apply it to the profile.</p>

MMS Scanning section of the New MMS Profile page

MMS Content Checksum	Select to add MMS Content Checksum in this MMS profile. Select the MMS content checksum list to apply it to the profile.
Pass Fragmented Messages	Select to pass fragmented MM3 and MM4 messages. Fragmented MMS messages cannot be scanned for viruses. If you do not select these options, fragmented MM3 and MM4 message are blocked.
Comfort Clients	<p>Select client comforting for MM1 and MM7 sessions.</p> <p>Since MM1 and MM7 messages use HTTP, MM1 and MM7 client comforting operates like HTTP client comforting.</p>
Comfort Servers	<p>Select server comforting for each protocol.</p> <p>Similar to client comforting, you can use server comforting to prevent server connection timeouts that can occur while waiting for the unit to buffer and scan large POST requests from slow clients.</p>
Interval (1-900 seconds)	Enter the time in seconds before client and server comforting starts after the download has begun, and the time between sending subsequent data.
Amount (1-10240 bytes)	The number of bytes sent by client or server comforting at each interval.
Oversized MMS Message	<p>Select Block or Pass for files and email messages exceeding configured thresholds for each protocol.</p> <p>The oversize threshold refers to the final size of the message, including attachments, after encoding by the client. Clients can use a variety of encoding types; some result in larger file sizes than the original attachment. As a result, a file may be blocked or logged as oversized even if the attachment is several megabytes smaller than the oversize threshold.</p>
Threshold (1KB - 800 MB)	Enter the oversized file threshold and select KB or MB. If a file is larger than the threshold the file is passed or blocked depending on the Oversized MMS Message setting. The web-based manager displays the allowed threshold range. The threshold maximum is 10% of the unit's RAM.

MMS Bulk Anti-Spam Detection options

You can use the MMS bulk email filtering options to detect and filter MM1 and MM4 message floods and duplicate messages. You can configure three thresholds that define a flood of message activity and three thresholds that define excessive duplicate messages. The configuration of each threshold includes the response actions for the threshold.

The configurable thresholds for each of the flood and duplicate sensors and must be enabled in sequence. For example, you can enable Flood Threshold 1 and Flood Threshold 2, but you cannot disable Flood Threshold 1 and enable Flood Threshold 2.

You can also add MSISDN to the bulk email filtering configuration and select a subset of the bulk email filtering options to applied to these individual MSISDNs.

You must first select MM1 and/or MM4 to detect excessive message duplicates. If excessive message duplicates are detected, the unit will perform the **Duplicate Message Action** for the specified duration.

You can configure three duplicate message thresholds and enable them with separate values and actions. They are labeled Duplicate Threshold 1 through 3 and must be enabled in sequence. For example, you can enable Duplicate Threshold 1 and Duplicate Threshold 2, but you cannot disable Duplicate Threshold 1 and enable Duplicate Threshold 2.

When traffic accepted by a security policy that contains an MMS profile with duplicate message configured receives MM1 or MM4 duplicate messages that match a threshold configured in the MMS protection profile, the unit performs the duplicate message action configured for the matching threshold.

You can configure three message flood thresholds and enable them with separate values and actions. They are labeled Flood Threshold 1 through 3 and must be enabled in sequence. For example, you can enable Flood Threshold 1 and Flood Threshold 2, but you cannot disable Flood Threshold 1 and enable Flood Threshold 2.

When traffic accepted by a security policy that contains an MMS protection profile with message flooding configured experiences MM1 or MM4 message flooding that matches a threshold configured in the MMS profile, the unit performs the message flood action configured for the matching threshold.

MMS Bulk Anti-Spam Detection

This section of the New MMS Profile page contains numerous sections where you can configure specific settings for flood threshold, duplicate threshold and recipient MSISDNs.

Message Flood

The message flood settings for each flood threshold. Expand each to configure settings for a threshold.

Flood Threshold 1	Expand to reveal the flood threshold settings for Flood Threshold 1. The settings for Flood Threshold 1 are the same for Flood Threshold 2 and 3.
Enable	Select to apply Flood Threshold 1 to the MSISDN exception.
Message Flood Window	Enter the period of time during which a message flood will be detected if the Message Flood Limit is exceeded. The message flood window can be 1 to 2880 minutes (48 hours).
Message Flood Limit	Enter the number of messages which signifies a message flood if exceeded within the Message Flood Window .
Message Flood Block Time	Enter the amount of time during which the unit performs the Message Flood Action after a message flood is detected.

Message Flood Action	Select one or more actions that the unit is to perform when a message flood is detected.
Flood Threshold 2	Expand to configure settings for Flood Threshold 2 or 3 respectively.
Flood Threshold 3	
Duplicate Message	
The duplicate message threshold settings. Expand each to configure settings for a threshold.	
MM1 Retrieve Duplicate Enable	Select to scan MM1 <code>mm1-retr</code> messages for duplicates. By default, <code>mm1-retr</code> messages are not scanned for duplicates as they may often be the same without necessarily being bulk or spam.
Enable	Select to enable the selected duplicate message threshold and to make the rest of the options available for configuration.
Duplicate Message Window	Enter the period of time during which excessive message duplicates will be detected if the Duplicate message Limit it exceeded. The duplicate message window can be 1 to 2880 minutes (48 hours).
Duplicate Message Limit	Enter the number of messages which signifies excessive message duplicates if exceeded within the Duplicate Message Window.
Duplicate Message Block Time	Enter the amount of time during which the unit will perform the Duplicate Message Action after a message flood is detected.
Duplicate Message Action	Select one or more actions that the unit is to perform when excessive message duplication is detected.
Duplicate Threshold 2	Expand to configure settings for Duplicate Threshold 2 or 3 respectively.
Duplicate Threshold 3	
Recipient MSISDN	
The recipient Mobile Subscriber Integrated Services Digital Network Number (MSISDN) settings for each recipient MSISDN. When you select Create New , you are automatically redirected to the New MSISDN page.	
You need to save the profile before you can add MSISDNs.	
Recipient MSISDN	The recipient MSISDN.
Flood Threshold 1	Check to enable Flood Threshold 1 settings for this MSISDN.
Flood Threshold 2	Check to enable Flood Threshold 2 settings for this MSISDN.
Flood Threshold 3	Check to enable Flood Threshold 3 settings for this MSISDN..

Duplicate Threshold 1	Check to enable Duplicate Threshold 1 settings for this MSISDN.
Duplicate Threshold 2	Check to enable Duplicate Threshold 2 settings for this MSISDN..
Duplicate Threshold 3	Check to enable Duplicate Threshold 3 settings for this MSISDN..
Edit	Modifies the settings of a Recipient MSISDN in the Recipient MSISDN list. When you select Edit , you are automatically redirected to the New MSISDN page.
Delete	Removes a Recipient MSISDN in the Recipient MSISDN list within the Recipient MSISDN section of the page.
New MSISDN page	
Create New	Creates a new Recipient MSISDN. When you select Create New , you are automatically redirected to the New MSISDN page.
Recipient MSISDN	Enter a name for the recipient MSISDN.
Flood Threshold 1	Select to apply Flood Threshold 1 to the MSISDN exception.
Flood Threshold 2	Select to apply Flood Threshold 2 to the MSISDN exception.
Flood Threshold 3	Select to apply Flood Threshold 3 to the MSISDN exception.
Duplicate Threshold 1	Select to apply Duplicate Threshold 1 to the MSISDN exception.
Duplicate Threshold 2	Select to apply Duplicate Threshold 2 to the MSISDN exception.
Duplicate Threshold 3	Select to apply Duplicate Threshold 3 to the MSISDN exception.

MMS Address Translation options

The sender's carrier endpoint is used to provide logging and reporting details to the mobile operator and to identify the sender of infected content.

When MMS messages are transmitted, the **From** field may or may not contain the sender's address. When the address is not included, the sender information will not be present in the logs and the unit will not be able to notify the user if the message is blocked unless the sender's address is made available elsewhere in the request.

The unit can extract the sender's address from an extended HTTP header field in the HTTP request. This field must be added to the HTTP request before it is received by the unit. If this field is present, it will be used instead of the sender's address in the MMS message for logging and notification. If this header field is present when a message is retrieved, it will be used instead of the **To** address in the message. If this header field is not present the content of the **To** header field is used instead.

Alternatively, the unit can extract the sender's address from a cookie.

You can configure MMS address translation to extract the sender's carrier endpoint so that it can be added to log and notification messages. You can configure MMS address translation settings to extract carrier endpoints from HTTP header fields or from cookies. You can also configure MMS address translation to add an endpoint prefix to

the extracted carrier endpoints. For more information, see Dynamic Profiles and Endpoints in the Authentication guide.

MMS Address Translation

Sender Address Source

Select to extract the sender's address from the **HTTP Header Field** or a **Cookie**. You must also specify the identifier that contains the carrier endpoint.

Enter the sender address identifier that includes the carrier endpoint. The default identifier is `x-up-calling-line-id`.

If the **Sender Address Source** is **HTTP Header Field**, the address and its identifier in the HTTP request header takes the format:

```
<Sender Address Identifier>: <MSISDN_value>
```

Where the `<MSISDN_value>` is the carrier endpoint. For example, the HTTP header might contain:

```
x-up-calling-line-id: 6044301297
```

where `x-up-calling-line-id` would be the Sender Address Identifier.

Sender Address Identifier

If the **Sender Address Source** is **Cookie**, the address and its identifier in the HTTP request header's `Cookie` field takes the format of attribute-value pairs:

```
Cookie: id=<cookie-id>;
```

```
<Sender Address Identifier>=<MSISDN Value>
```

For example, the HTTP request headers might contain:

```
Cookie: id=0123jff!a;x-up-calling-line-id=6044301297
```

where `x-up-calling-line-id` would be the **Sender Address Identifier**.

Convert Sender Address From / To HEX

Select to convert the sender address from ASCII to hexadecimal or from hexadecimal to ASCII. This is required by some applications.

Add Carrier Endpoint Prefix for Logging / Notification

Select the following to enable adding endpoint prefixes for logging and notification.

Enable

Select to enable adding the country code to the extracted carrier endpoint, such as the MSISDN, for logging and notification purposes. You can limit the number length for the test numbers used for internal monitoring without a country code.

MMS Address Translation

Prefix	Enter a carrier endpoint prefix to be added to all carrier endpoints. Use the prefix to add extra information to the carrier endpoint in the log entry.
Minimum Length	Enter the minimum length of the country code information being added. If this and Maximum Length are set to zero (0), length is not limited.
Maximum Length	Enter the maximum length of the country code information being added. If this and Minimum Length are set to zero (0), length is not limited.

MMS Notifications

MMS notifications are messages that a unit sends when an MMS profile matches content in an MM1, MM3, MM4 or MM7 session. For example, the MMS profile detects a virus or uses content blocking to block a web page, text message or email. You can send notifications to the sender of the message using same protocol and the addressing headers in the original message. You can also configure MMS notifications to send notification messages to another destination (such as a system administrator) using the MM1, MM3, MM4 or MM7 protocol.

You need to enable one or more **Notification Types** or you can add an **Antivirus Notification List** to enable sending notifications.

You can also use MMS notifications options to configure how often notifications are sent. The unit sends notification messages immediately for the first event, then at a configurable interval if events continue to occur. If the interval does not coincide with the window of time during which notices may be sent, the unit waits to send the notice in the next available window. Subsequent notices contain a count of the number of events that have occurred since the previous notification.

There are separate notifications for each notification type, including virus events. Virus event notifications include the virus name. Up to three viruses are tracked for each user at a time. If a fourth virus is found, one of the existing tracked viruses is removed from the list.

The notifications are MM1 `m-send-req` messages sent from the unit directly to the MMSC for delivery to the client. The host name of the MMSC, the URL to which `m-send-req` messages are sent, and the port must be specified.

MMS Notification

Antivirus Notification List	<p>Optionally select an antivirus notification list to select a list of virus names to send notifications for. The unit sends a notification message whenever a virus name or prefix in the antivirus notification list matches the name of a virus detected in a session scanned by the MMS protection profile. Select Disabled if you do not want to use a notification list.</p> <p>Instead of selecting a notification list you can configure the Virus ScanNotification Type to send notifications for all viruses.</p>
------------------------------------	--

MMS Notification

Message Protocol	<p>In each column, select the protocol used to send notification messages. You can use a different protocol to send the notification message than the protocol on which the violation was sent. The MMS Notifications options change depending on the message protocol that you select.</p> <p>If you select a different message protocol, you must also enter the User Domain. If selecting MM7 you must also enter the Message Type.</p>
Message Type	Select the MM7 message type to use if sending notifications using MM7. Options include deliver.REQ and submit.REQ
Detect Server Details	<p>Select to use the information in the headers of the original message to set the address of the notification message. If you do not select this option, you can enter the required addressing information manually.</p> <p>You cannot select Detect Server Details if you are sending notification messages using a different message protocol.</p> <p>If you select Detect Server Details, you cannot change the Port where the notification is being sent.</p>
Hostname	Enter the FQDN or the IP address of the server where the notifications will be sent.
URL	<p>Enter the URL of the server. For example if the notification is going to www.example.com/home/alerts , the URL is /home/alerts.</p> <p>This option is available only when Message Protocol is mm1 or mm7.</p>
Port	<p>Enter the port number of the server.</p> <p>You cannot change the Port if Detect Server Details is enabled.</p>
Username	<p>Enter the user name required for sending messages using this server (optional).</p> <p>This option is available only when Message Protocol is mm7.</p>
Password	<p>Enter the password required for sending messages using this server (optional).</p> <p>This option is available only when Message Protocol is mm7.</p>

MMS Notification	
VASP ID	<p>Enter the value-added-service-provider (VASP) ID to be used when sending a notification message. If a VAS is not offered by the mobile provider, it is offered by a third party or a VAS provider or content provider (CP).</p> <p>This option is available only when Message Protocol is mm7.</p>
VAS ID	<p>Enter the value-added-service (VAS) ID to be used when sending a notification message. A VAS is generally any service beyond voice calls and fax.</p> <p>This option is available only when Message Protocol is mm7.</p>
All Notification Types	<p>In each column, select notification for all MMS event types for that MMS protocol, then enter the amount of time and select the time unit for notice intervals.</p> <p>Alternatively, expand All Notification Types, and then select notification for individual MMS event types for each MMS protocol. Then enter the amount of time and select the time unit for notice intervals.</p> <p>Not all event types are available for all MMS protocols.</p>
Content Filter	In each column, select to notify when messages are blocked by the content filter, then enter the amount of time and select the time unit for notice intervals.
File Block	In each column, select to notify when messages are blocked by file block, then enter the amount of time and select the time unit for notice intervals.
Carrier Endpoint Block	In each column, select to notify when messages are blocked, then enter the amount of time and select the time unit for notice intervals.
Flood	In each column, select to notify when message flood events occur, then enter the amount of time and select the time unit for notice intervals.
Duplicate	In each column, select to notify when duplicate message events occur, then enter the amount of time and select the time unit for notice intervals.
MMS Content Checksum	In each column, select to notify when the content within an MMS message is scanned and banned because of the checksum value that was matched.
Virus Scan	In each column, select to notify when the content within an MMS message is scanned for viruses.

MMS Notification	
Notifications Per Second Limit	For each MMS protocol, enter the number of notifications to send per second. If you enter zero (0) , the notification rate is not limited.
Day of Week	For each MMS protocol, select the days of the week the unit is allowed to send notifications.
Window Start Time	For each MMS protocol, select the time of day to begin the message alert window. By default, the message window starts at 00:00. You can change this if you want to start the message window later in the day. When configured, notification outside this window will not be sent.
Window Duration	For each MMS protocol, select the time of day at which to end the message alert window. By default, the message window ends at 00:24. You can change this if you want to end the message window earlier in the day. When configured, notification outside this window will not be sent

DLP Archive options

Select DLP archive options to archive MM1, MM3, MM4, and MM7 sessions. In addition to the MMS profile's DLP archive options, you can:

- Archive MM1 and MM7 message floods
- Archive MM1 and MM7 duplicate messages
- Select **DLP archiving** for carrier endpoint patterns in a **Carrier Endpoint List** and select the **Carrier Endpoint Block** option in the **MMS Scanning** section of an MMS Profile

The unit only allows one sixteenth of its memory for transferring content archive files. For example, for units with 128 MB RAM, only 8 MB of memory is used when transferring content archive files. Best practices dictate to not enable full content archiving if antivirus scanning is also configured because of these memory constraints.

DLP Archive	
Display DLP meta-information on the system dashboard	Select each required protocol to display the content archive summary in the Log and Archive Statistics dashboard widget on the System Dashboard.

DLP Archive

Archive to FortiAnalyzer/FortiGuard	Select the type of archiving that you want for the protocol (MM1, MM3, MM4, and MM7). You can choose from Full, Summary or None.
	None — Do not send content archives.
	Summary — Send content archive metadata only. Includes information such as date and time, source and destination, request and response size, and scan result.
	<p>Full — Send content archive both metadata and copies of files or messages.</p> <p>In some cases, FortiOS Carrier may not archive content, or may make only a partial content archive, regardless of your selected option. This behavior varies by prerequisites for each protocol.</p> <p>This option is available only if a FortiAnalyzer unit or FortiGuard Analysis and Management Service is configured.</p>

Logging

You can enable logging in an MMS profile to write event log messages when the MMS profile options that you have enabled perform an action. For example, if you enable MMS antivirus protection, you could also use the MMS profile logging options to write an event log message every time a virus is detected.

You must first configure how the unit stores log messages so that you can then record these logs messages. For more information, see the FortiOS Handbook Logging and Reporting guide.

Logging

MMS-Antivirus	If antivirus settings are enabled for this MMS profile, select the following options to record Antivirus Log messages.
Viruses	Record a log message when this MMS profile detects a virus.
Blocked Files	Record a log message when antivirus file filtering enabled in this MMS profile blocks a file.
Oversized Files/Emails	Record a log message when this MMS profile encounters an oversized file or email message. Oversized files and email messages cannot be scanned for viruses.
MMS Scanning	If MMS scanning settings are enabled for this MMS profile, select the following options to record Email Filter Log messages.
Notification Messages	Select to log the number of MMS notification messages sent.

Logging	
Bulk Messages	Select to log MMS Bulk AntiSpam events. You must also select which protocols to write log messages for in the MMS bulk email filtering part of the MMS profile.
Carrier Endpoint Filter Block	Select to log MMS carrier endpoint filter events, such as MSISDN filtering.
MMS Content Checksum	Select to log MMS content checksum activity.
Content Block	Select to log content blocking events.

MMS Content Checksum

The MMS Content Checksum menu allows you to configure content checksum lists.

Configure MMS content checksum lists in **Security Profiles > MMS Content Checksum** using the following table.

MMS Content Checksum	
Lists each individual content checksum list that you created. On this page, you can edit, delete or create a content checksum list.	
Create New	Creates a new MMS content checksum list. When you select Create New , you are automatically redirected to the New List. This page provides a name field and comment field. You must enter a name to go to MMS Content Checksum Settings page.
Edit	Modifies settings to a MMS content checksum. When you select Edit , you are automatically redirected to the MMS Content Checksum Settings page.
Delete	Removes an MMS content checksum from the page.
	To remove multiple content checksum lists from within the list, on the MMS Content Checksum page, in each of the rows of the content checksum lists you want removed, select the check box and then select Delete .
	To remove all content checksum lists from list, on the MMS Content Checksum page, select the check box in the check box column and then select Delete .
Name	The name of the MMS content checksum list that you created.
# Entries	The number of checksums that are included in the content checksum list.

MMS Profiles	The MMS profile or profiles that have the MMS content checksum list applied. For example if two different MMS profiles use this content checksum list, they will both be listed here.
Comments	A description given to the MMS content checksum.
Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (Security Profiles > AntiVirus > Profiles), 1 appears in Ref. .</p> <p>To view the location of the referenced object, select the number in Ref., and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.

Notification List

The Notification List menu allows you to configure a list of viruses. This virus list provides a list for scanning viruses in MMS messages. You can use one virus list in multiple MMS profiles, and configure multiple virus lists.

Notification list configuration settings

The following are notification list configuration settings in **Security Profiles > Notification List**.

Notification List	Lists all the notification lists that you created. On this page you can edit, delete or create a new notification list.
Create New	Creates a new notification list. When you select Create New , you are automatically redirected to the New List page. You must enter a name to go to the Notification List Settings page.

Edit	Modifies settings within the notification list. When you select Edit , you are automatically redirected to the Notification List Settings page.
Delete	<p>Removes a notification list from the list on the Notification List page.</p> <p>To remove multiple notification lists from within the list, on the Notification List page, in each of the rows of the notification lists you want removed, select the check box and then select Delete.</p> <p>To remove all notification lists from the list, on the Notification List page, select the check box in the check box column and then select Delete.</p>
Name	The name of the MMS content checksum list that you created.
# Entries	The number of checksums that are included in that content checksum list.
MMS Profiles	The MMS profile or profiles that are associated with
Comments	A description given to the MMS notification list.
Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (Security Profiles > Antivirus > Profiles), 1 appears in Ref. .</p> <p>To view the location of the referenced object, select the number in Ref., and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.
Notification List Settings	
Provides settings for configuring a notification list, which is a list of viruses and is used for scanning viruses in MMS messages. This list is called the Antivirus Notification List in an MMS profile.	

Name	If editing the name of a notification list, enter the new name in this field. You must select OK to save the change.
Comments	If you want to enter a comment, enter the comment in the field. You must select OK to save the change.
Create New	Creates a notification entry in the list. When you select Create New , you are automatically redirected to the New Entry page.
Edit	Modifies settings within a notification list. When you select Edit , you are automatically redirected to the Edit Entry page.
Delete	Removes a notification entry from the list on the page. To remove multiple notification entries from within the list, on the Notification List Settings page, in each of the rows of the entries you want removed, select the check box and then select Delete . To remove all notification entries from the list, on the Notification List Settings page, select the check box in the check box column and then select Delete .
Enable	Enables a notification entry that is disabled.
Disable	Disables a notification entry so that it is not active and available for use, but it is not deleted.
Remove All Entries	Removes all notification entries that are listed on the Notification List Settings page.
Enable	Displays whether or not the checksum is enabled.
Virus Name/Profile	The name of the virus that was added to the list.
Entry Type	The type of match that will be used to match the virus stated in the notification list to the actual virus that is found.
New Entry page	
Virus Name/Profile	Enter the virus name.
Entry Type	Select the type of match that will be used to match the virus stated in the notification list to the actual virus that is found.
Enable	Select to enable the virus in the list.

Message Flood

The convenience offered by MM1 and MM4 messaging can be abused by users sending spam or attempting to overload the network with an excess of messages. MMS flood prevention can help prevent this type of abuse. A

message flood occurs when a single subscriber sends a volume of messages that exceed the flood threshold that you set. The threshold defines the maximum number of messages allowed, the period during which the subscriber sent messages are considered, and the length of time the sender is restricted from sending messages after a flood is detected. For example, for the first threshold you may determine that any subscriber who sends more than 100 MM1 messages in an hour (60 minutes) will have all outgoing messages blocked for 30 minutes.

Action	Description
Log	Add a log entry indicating that a message flood has occurred. You must also enable logging by going to Security Profiles > MMS Profile , <applicable profile> > Logging > MMS Scanning > Bulk Messages , and toggling on the checkbox.
DLP Archive	Save the first message to exceed the flood threshold, or all the messages that exceed the flood threshold, in the DLP archive. DLP archiving flood messages may not always produce useful results. Since different messages can be causing the flood, reviewing the archived messages may not be a good indication of what is causing the problem since the messages could be completely random.
All messages	All the messages that exceed the flood threshold will be saved in the DLP archive.
First message only	Save only the first message to exceed the flood threshold in the DLP archive. Other messages in the flood are not saved. For message floods this may not produce much useful information since a legitimate message could trigger the flood threshold.
Intercept	Messages that exceed the flood threshold are passed to the recipients, but if quarantine is enabled for intercepted messages, a copy of each message will also be quarantined for later examination. If the quarantine of intercepted messages is disabled, the Intercept action has no effect.
Block	Messages that exceed the flood threshold are blocked and will not be delivered to the message recipients. If quarantine is enabled for blocked messages, a copy of each message will be quarantined for later examination.
Alert Notification	<p>If the flood threshold is exceeded, the Carrier-enabled FortiGate unit will send an MMS flood notification message.</p> <p>In the web-based manager when Alert Notification is selected it displays the fields to configure the notification.</p>

Flood protection for MM1 messages prevents your subscribers from sending too many messages to your MMSC. Configuring flood protection for MM4 messages prevents another service provider from sending too many messages from the same subscriber to your MMSC.

Message flood configuration settings

The following are message flood configuration settings in **Security Profiles > Message Flood**.

Message Flood	
Lists the large amount of messages that are being sent to you from outside sources.	
Delete	<p>Removes messages from the list.</p> <p>To remove multiple messages from within the list, on the Message Flood page, in each row of the messages you want removed, select the check box and then select Delete.</p> <p>To remove all messages from the list, on the Message Flood page, select the check box in the check box column and then select Delete.</p>
Remove All Entries	Removes all messages from the list.
Protocol	Sorts/filters by the protocol used.
MMS Profile	Sorts/filters by the MMS profile that is used.
Sender	Sorts/filters by the sender's email address.
Level	Sorts/filters by the level of severity of the message.
Count	The count column can be up or down and these settings can be turned off by selecting beside the column's name.
Window Size (minutes)	The time in minutes.
Timer (minutes:seconds)	The time in seconds and in minutes. The timer column can be up or down and these settings turned off by selecting beside the column's name.
Page Controls	Use to navigate through the list.

Duplicate Message

Duplicate message protection for MM1 messages prevents multiple subscribers from sending duplicate messages to your MMSC. Duplicate message protection for MM4 messages prevents another service provider from sending duplicate messages from the same subscriber to your MMSC.

The unit keeps track of the sent messages. If the same message appears more often than the threshold value that you have configured, action is taken. Possible actions are logging the duplicate messages, blocking or intercepting them, archiving, and sending an alert to inform an administrator that duplicate messages are occurring.

Duplicate message configuration settings

View duplicate messages in **Security Profiles > Duplicate Message**.

Duplicate Message

Lists duplicates of messages that were sent to you.

Delete	Removes a message from the list.
	To remove multiple duplicate messages from within the list, on the Message Flood page, in each row of the messages you want removed, select the check box and then select Delete .
	To remove all duplicate messages from the list, on the Message Flood page, select the check box in the check box column and then select Delete .
Page Controls	Use to navigate through the list.
Remove All Entries	Removes all duplicate messages from the list.
Protocol	Sorts/filters by the protocol used.
MMS Profile	Sorts/filters by the MMS profile that logs the detection.
Checksum	Sorts/filters by the checksum of the MMS message.
Level	Sorts/filters by the level of severity of the message.
Count	Displays the number of messages in the last window of time.
Window Size (minutes)	The period of time during which a message flood will be detected if the Message Flood Limit is exceeded.
Timer (minutes:seconds)	Either the time left in the window if the message is unflagged, or the time until the message will be unflagged if it is already flagged.

Carrier Endpoint Filter Lists

A carrier endpoint filter list contains carrier endpoint patterns. A pattern can match one carrier endpoint or can use wildcards or regular expressions to match multiple carrier endpoints. For each pattern, you select the action that the unit takes on a message when the pattern matches a carrier endpoint in the message. Actions include blocking the message, exempting the message from MMS scanning, and exempting the message from all scanning. You can also configure the pattern to intercept the message and content archive the message to a FortiAnalyzer unit.

Carrier endpoint filter lists configuration settings

The following are Carrier endpoint filter list configuration settings in **Security Profiles > Carrier Endpoint Filter Lists**.

Carrier Endpoint Filter Lists

Lists all the endpoint filters that you created. On this page, you can edit, delete or create a new endpoint filter list.

Create New	Creates a new endpoint filter list. When you select Create New , you are automatically redirected to the New List page. You must enter a name to go to the Carrier Endpoint Filter Lists Settings page.
Edit	Modifies settings within an endpoint filter list in the list.
Delete	<p>Removes an endpoint filter in the list.</p> <p>To remove multiple endpoint filter lists from within the list, on the Carrier Endpoint Filter List page, in each of the rows of the endpoint filter lists you want removed, select the check box and then select Delete.</p> <p>To remove all endpoint filter lists from the list, on the Carrier Endpoint Filter List page, select the check box in the check box column and then select Delete.</p>
Name	The name of the endpoint filter.
# Entries	The number of carrier endpoint patterns in each carrier endpoint filter list.
MMS Profiles	The MMS profile that the carrier endpoint filter list is added to.
Comments	A description about the endpoint filter.

Ref.	<p>Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (Security Profiles > Antivirus > Profiles), 1 appears in Ref. .</p> <p>To view the location of the referenced object, select the number in Ref., and the Object Usage window appears displaying the various locations of the referenced object.</p> <p>To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:</p> <ul style="list-style-type: none"> • View the list page for these objects – automatically redirects you to the list page where the object is referenced at. • Edit this object – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page. • View the details for this object – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.
-------------	---

Carrier Endpoint Filter Lists Settings

Provides settings for configuring an endpoint filter.

Name	The name you entered on the New List page, after selecting Create New on the Carrier Endpoint Filter page.
Comments	A description about the endpoint filter. You can add one here if you did not enter one on the New List page.
Create New	Creates a new endpoint filter list. When you select Create New , you are automatically redirected to the New Entry page.
Edit	Select to modify the settings of a pattern in the list.
Delete	Select to remove a pattern in the list.
Enable	Enables a disabled pattern in the list.
Disable	Disables a pattern in the list.
Remove All Entries	Removes all patterns in the list on the Carrier Endpoint Filter Lists Settings page.
Enable	Indicates whether or not the pattern is enabled.

Pattern	Enter or change the pattern that FortiOS Carrier uses to match with carrier endpoints. The pattern can be a single carrier endpoint or consist of wildcards or Perl regular expressions that will match more than one carrier endpoint. Set Pattern Type to correspond to the pattern that you want to use.
Action	Select the action taken by FortiOS Carrier for messages from a carrier endpoint that matches the carrier endpoint pattern:
Pattern Type	The type of pattern chosen.
New Entry page	
Pattern	Enter or change the pattern that FortiOS Carrier uses to match with carrier endpoints. The pattern can be a single carrier endpoint or consist of wildcards or Perl regular expressions that will match more than one carrier endpoint. Set Pattern Type to correspond to the pattern that you want to use.
Action(s)	<p>Select the action taken by FortiOS Carrier for messages from a carrier endpoint that matches the carrier endpoint pattern:</p> <p>Action(s) can be:</p> <ul style="list-style-type: none"> • None • Block • Exempt from mass MMS • Exempt from all scanning
Content Archive	MMS messages from the carrier endpoint are delivered, the message content is DLP archived according to MMS DLP archive settings. Content archiving is also called DLP archiving.
Pattern Type	<p>Select a pattern type as one of Single Carrier Endpoint, Wildcard or Regular Expression.</p> <p>Wildcard and Regular Expression will match multiple patterns where Single Carrier Endpoint matches only one.</p>
Enable	Select to enable this carrier endpoint filter pattern.

Message flood protection

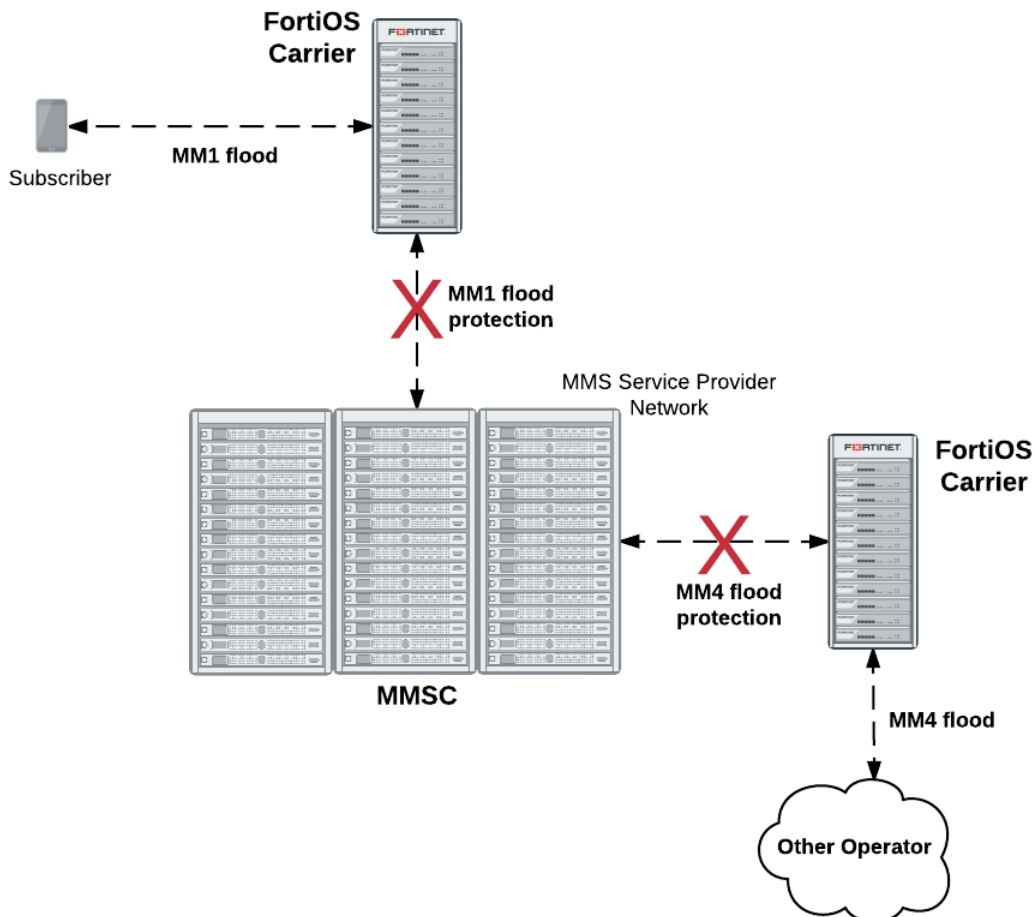
The convenience offered by MM1 and MM4 messaging can be abused by users sending spam or attempting to overload the network with an excess of messages. MMS flood prevention can help prevent this type of abuse.

Overview

Flood protection for MM1 messages prevents your subscribers from sending too many messages to your MMSC. Configuring flood protection for MM4 messages prevents another service provider from sending too many

messages from the same subscriber to your MMSC.

MM1 and MM4 flood protection



The FortiOS Carrier unit keeps track of the number of messages each subscriber sends for the length of time you specify. If the number of messages a subscriber sends exceeds the threshold, a configured action is taken. Possible actions are logging the flood, blocking or intercepting messages in the flood, archiving the flood messages, and sending an alert message to inform the administrator that the flood is occurring.

You can create three different thresholds to take different levels of action at different levels of activity.

With this highly configurable system, you can prevent subscribers from sending more messages than you determine is acceptable, or monitor anyone who exceeds the thresholds.

Setting message flood thresholds

A message flood occurs when a single subscriber sends a volume of messages that exceeds the flood threshold you set. The threshold defines the maximum number of messages allowed, the period during which the subscriber sent messages are considered, and the length of time the sender is restricted from sending messages after a flood is detected.

If a subscriber exceeds the message flood threshold and is blocked from sending more messages, any further attempts to send messages will re-start the block period. You must also enable logging for **MMS Scanning > Bulk Messages** in the Logging section of the MMS protection profile.



A subscriber is still able to receive messages while they are blocked from sending messages.

Example

For example, for the first threshold you may determine that any subscriber who sends more than 100 MM1 messages in an hour (60 minutes) will have all messages blocked for half an hour (30 minutes).

Using this example, if the subscriber exceeds the flood threshold, they are blocked from sending message for 30 minutes. If the subscriber tries to send any message after 15 minutes, the message will be blocked and the block period will be reset again to 30 minutes. The block period must expire with no attempts to send a message. Only then will the subscriber be allowed to send more messages.

To configure MM1 message flood threshold - web-based manager

1. Go to **Security Profiles > MMS Profile**.
2. Select **Create New**.
3. Enter `MM1 flood` for **Profile Name**.
4. Expand **MMS Bulk Email Filtering Detection**.
5. Enter the following information, and select **OK**.

MM1 (first column)	
Enable	Enable
Message Flood Window	60 minutes
Message Flood Limit	100
Message Flood Block Time	30 minutes
Message Flood Action	Block

To configure MM1 message flood threshold - CLI

```
config firewall mms-profile
edit profile_name
config flood mm1
set status1 enable
set window1 60
set limit1 100
set action1 block
set block-time1 30
end
end
```

The threshold values that you set for your network will depend on factors such as how busy your network is and the kinds of problems that your network and your subscribers encounter. For example, if your network is not too busy you may want to set message flood thresholds relatively high so that only an exceptional situation will exceed a flood threshold. Then you can use log messages and archived MMS messages to determine what caused the flood.

If your subscribers are experiencing problems with viruses that send excessive amounts of messages, you may want to set thresholds lower and enable blocking to catch problems as quickly as possible and block access to keep the problem from spreading.

Flood actions

When the Carrier-enabled FortiGate unit detects a message flood, it can take any combination of the five actions that you can configure for the flood threshold. For detailed options, see Message Flood.

Notifying administrators of floods

You can configure alert notifications for message floods by selecting the Alert Notification message flood action.

The FortiOS Carrier unit sends alert notifications to administrators using the MM1, MM3, MM4, or MM7 content interface. To send an alert notification you must configure addresses and other settings required for the content interface.

For example, to send notifications using the MM1 content interface you must configure a source MSISDN, hostname, URL, and port to which to send the notification. You can also configure schedules for when to send the notifications.

Finally you can add multiple MSISDN numbers to the MMS protection profile and set which flood thresholds to send to each MSISDN.

Example — three flood threshold levels with different actions for each threshold

You can set up to three threshold levels to take different actions at different levels of activity.

The first example threshold records log messages when a subscriber's handset displays erratic behavior by sending multiple messages using MM1 at a relatively low threshold. The erratic behavior could indicate a problem with the subscriber's handset. For example, you may have determined for your network that if a subscriber sends more the 45 messages in 30 minutes that you want to record log messages as a possible indication or erratic behavior.

From the web-based manager in an MMS profile set message **Flood Threshold 1** to:

Enable	Selected
Message Flood Window	30 minutes
Message Flood Limit	45
Message Flood Action	Log

From the CLI:

```
config firewall mms-profile
  edit profile_name
    config flood mm1
```

```

        set status1 enable
        set window1 30
        set limit1 45
        set action1 log
    end
end

```

Set a second higher threshold to take additional actions when a subscriber sends more that 100 messages in 30 minutes. Set the actions for this threshold to log the flood, archive the message that triggered the second threshold, and block the sender for 15 minutes.

From the web-based manager in an MMS profile set message **Flood Threshold 2** to:

Enable	Selected
Message Flood Window	30 minutes
Message Flood Limit	100
Message Block Time	15 minutes
Message Flood Action	Log, DLP archive First message only, Block

From the CLI:

```

config firewall mms-profile
  edit profile_name
    config flood mm1
      set status2 enable
      set window2 30
      set limit2 100
      set action2 block log archive-first
      set block-time2 15
    end
  end
end

```

Set the third and highest threshold to block the subscriber for an extended period and sand an administrator alert if the subscriber sends more than 200 messages in 30 minutes. Set the actions for this threshold to block the sender for four hours (240 minutes), log the flood, archive the message that triggered the third threshold, and send an alert to the administrator.

From the web-based manager in an MMS profile set message **Flood Threshold 3** to:

Enable	Selected
Message Flood Window	30 minutes
Message Flood Limit	200
Message Block Time	240 minutes
Message Flood Action	Log, Block, Alert Notification

Because you have selected the **Alert Notification** action you must also configure alert notification settings. For this example, the source MSISDN is 5551234—telephone number 555-1234. When administrators receive MMS messages from this MSISDN they can assume a message flood has been detected.

In this example, alert notifications are sent by the FortiOS Carrier unit to the MMSC using MM1. The host name of the MMSC is `mmscexample`, the MMSC URL is `/`, and the port used by the MMSC is 80. In this example, the alert notification window starts at 8:00am and extends for eight hours on weekdays (Monday-Friday) and the minimum interval between message flood notifications is two hours.

Source MSISDN	5551234
Message Protocol	MM1
Hostname	mmscexample
URL	/
Port	80
Notifications Per Second Limit	0
Window Start Time	8:00
Window Duration	8:00
Day of Week	Mon, Tue, Wed, Thu, Fri, Sat
Interval	2 hours

From the CLI:

```
config firewall mms-profile
  edit profile_name
    config notification alert-flood-1
      set alert-src-msisdn 5551234
      set msg-protocol mm1
      set mmsc-hostname mmscexample
      set mmsc-url /
      set mmsc-port 80
      set rate-limit 0
      set tod-window-start 8:00
      set tod-window-duration 8:00
      set days-allowed monday tuesday wednesday thursday friday
      set alert-int 2
      set alert-int-mode hours
    end
```

You must also add the MSISDNs of the administrators to be notified of the message flood. In this example, the administrator flood threshold 3 alert notifications are sent to one administrator with MSISDN 5554321.

To add administrator's MSISDNs for flood threshold 3 from the web-based manager when configuring a protection profile, select **MMS Bulk Email Filtering Detection > Recipient MSISDN > Create New**.

MSISDN	5554321
Flood Level 3	Select

From the CLI:

```
config firewall mms-profile
  edit profile_name
    config notif-msisdn
      edit 5554321
        set threshold flood-thresh-3
      end
    end
  end
```

Notifying message flood senders and receivers

The FortiOS Carrier unit does not send notifications to the sender or receiver that cause a message flood. If the sender or receiver is an attacker and is explicitly informed that they have exceeded a message threshold, the attacker may try to determine the exact threshold value by trial and error and then find a way around flood protection. For this reason, no notification is set to the sender or receiver.

However, FortiOS Carrier does have replacement messages for sending reply confirmations to MM1 senders and receivers and for MM4 senders for blocked messages identified as message floods. For information about how FortiOS Carrier responds when message flood detection blocks a message, see and MMS duplicate messages and message floods.

Responses to MM1 senders and receivers

When the FortiOS Carrier unit identifies an MM1 message sent by a sender to an MMSC as a flood message and blocks it, the FortiOS Carrier unit returns a message submission confirmation (`m-send.conf`) to the sender — otherwise the sender's handset would keep retrying the message. The `m-send.conf` message is sent only when the MM1 message flood action is set to Block. For other message flood actions the message is actually delivered to the MMSC and the MMSC sends the `m-send.conf` message.

You can customize the `m-send.conf` message by editing the **MM1 send-conf flood message** MM1 replacement message (from the CLI the `mm1-send-conf-flood` replacement message). You can customize the response status and message text for this message. The default response status is "Content not accepted". To hide the fact that FortiOS Carrier is responding to a flood, you can change the response status to "Success". The default message text informs the sender that the message was blocked. You could change this to something more generic.

For example, the following command sets the submission confirmation response status to "Success" and changes the message text to "Message Sent OK":

```
config system replacemsg mm1 mm1-send-conf-flood
  set rsp-status ok
  set rsp-text "Message Sent OK"
end
```

When the FortiOS Carrier unit identifies an MM1 message received by a receiver from an MMSC as a flood message and blocks it, the FortiOS Carrier unit returns a message retrieval confirmation (`m-retrieve.conf`) to the sender (otherwise the sender's handset would keep retrying the message). The `m-retrieve.conf` message is sent only when the MM1 message flood action is set to Block. For other message flood actions the message is actually delivered to the receiver, so the MMSC sends the `m-retrieve.conf` message.

You can customize the m-retrieve.conf message by editing the **MM1 retrieve-conf flood message** MM1 replacement message (from the CLI the `mm1-retr-conf-flood` replacement message). You can customize the class, subject, and message text for this message.

For example, you could use the following command make the response more generic:

```
config system replacemsg mm1 mm1-retr-conf-flood
    set subject "Message blocked"
    set message "Message temporarily blocked by carrier"
end
```

Forward responses for MM4 message floods

When the FortiOS Carrier unit identifies an MM4 message as a flood message and blocks it, the FortiOS Carrier unit returns a message forward response (MM4_forward.res) to the forwarding MMSC (otherwise the forwarding MMSC would keep retrying the message). The MM4_forward.res message is sent only when the MM4 message flood action is set to Block and the MM4-forward.req message requested a response. For more information, see and MMS duplicate messages and message floods.

You can customize the MM4_forward.res message by editing the **MM4 flood message** MM4 replacement message (from the CLI the `mm4-flood` replacement message). You can customize the response status and message text for this message. The default response status is "Content not accepted" (`err-content-not-accept`). To hide the fact that the FortiOS Carrier unit is responding to a flood, you can change the response status to "Success". The default message text informs the sender that the message was blocked. You could change this to something more generic.

For example, the following command sets the submission confirmation response status to "Success" and changes the message text to "Message Sent OK" for the MM4 message forward response

```
config system replacemsg mm4 mm4-flood
    set rsp-status ok
    set rsp-text "Message Forwarded OK"
end
```

Viewing DLP archived messages

If **DLP Archive** is a selected message flood action, the messages that exceed the threshold are saved to the MMS DLP archive. The default behavior is to save all of the offending messages, but you can configure the DLP archive setting to save only the first message that exceeds the threshold. This still provides a sample of the offending messages without requiring as much storage.

To select only the first message in a flood for DLP archiving - web-based manager

1. Go to **Security Profiles > MMS Profile**.
2. Edit an existing MMS Profile.
3. Expand the **MMS Bulk Email Filtering Detection** section, the **Message Flood** subsection, and the desired **Flood Threshold** subsection.
4. Next to **DLP Archive**, select **First message only** from the drop down menu.
5. Select **OK**.

Order of operations: flood checking before duplicate checking

Although duplicate checking involves only examination and comparison of message contents and not the sender or recipient, and flood checking involves only totaling the number of messages sent by each subscriber regardless

of the message content, there are times when a selection of messages exceed both flood and duplicate thresholds.

The Carrier-enabled FortiGate unit checks for message floods before checking for duplicate messages. Flood checking is less resource-intensive and if the flood threshold invokes a **Block** action, the blocked messages are stopped before duplicate checking occurs. This saves both time and FortiOS Carrier system resources.



The duplicate scanner will only scan content. It will not scan headers. Content must be exactly the same. If there is any difference at all in the content, it will not be considered a duplicate.

Bypassing message flood protection based on user's carrier endpoints

You can use carrier endpoint filtering to exempt MMS sessions from message flood protection. Carrier endpoint filtering matches carrier endpoints in MMS sessions with carrier endpoint patterns.

If you add a carrier endpoint pattern to a filter list and set the action to exempt from mass MMS, all messages from matching carrier endpoints bypass message flood protection. This allows legitimate bulk messages, such as system outage notifications, to be delivered without triggering message flood protection.

For more information on carrier endpoints, see the User Authentication chapter of the FortiOS Handbook.

Configuring message flood detection

To have the Carrier-enabled FortiGate unit check for message floods, you must first configure the flood threshold in an MMS profile, select the MMS profile in a security policy. All the traffic examined by the security policy will be checked for message floods according to the threshold values you set in the MMS profile.

Configure the MMS profile - web-based manager

1. Go to **Firewall Objects > MMS Profile**.
2. If you are editing an MMS profile, select the **Edit** icon of the MMS profile.
If you are creating a new MMS profile, select **Create New** and enter a profile name.
3. Expand **MMS Bulk Email Filtering Detection**.
4. Expand **Message Flood**.
5. Expand **Flood Threshold 1**.
6. Select the **Enable** check box for MM1 messages, MM4 messages, or both.
7. In the **Message Flood Window** field, enter the length of time the Carrier-enabled FortiGate unit will keep track of the number of messages each subscriber sends.
If the Carrier-enabled FortiGate unit detects the quantity of messages specified in the **Message Flood Limit** sent during the number of minutes specified in the **Message Flood Window**, a message flood is in progress.
8. In the **Message Flood Limit** field, enter the number of messages required to trigger the flood.
9. In the **Message Flood Block Time** field, enter the length of time a user will be blocked from sending messages after causing the message flood.
10. Select the message flood actions the Carrier-enabled FortiGate unit will take when the message flood is detected.
11. Select **OK**.

Configure the security policy - web-based manager

1. Go to **Policy**.
2. Select the **Edit** icon of the security policy that controls the traffic in which you want to detect message floods.
3. Select the **MMS Profile** check box to enable the use of a protection profile.
4. Select the MMS protection profile from the list.
5. Select **OK**.

Sending administrator alert notifications

When message floods are detected, the Carrier-enabled FortiGate unit can be configured to notify you immediately with an MMS message. Enable this feature by selecting Alert Notification in the message flood action. Each message flood threshold can be configured separately.

Configuring how and when to send alert notifications

You can configure different alert notifications for MM1 and MM4 message floods. You can configure the FortiOS Carrier unit to send these alert notifications using the MM1, MM3, MM4, or MM7 content interface. Each of these content interfaces requires alert notification settings that the FortiOS Carrier unit uses to communicate with a server using the selected content interface.

For the MM1 content interface you require:

- The hostname of the server
- The URL of the server (usually "/")
- The server port (usually 80)

For the MM3 and MM4 content interfaces you require:

- The hostname of the server
- The server port (usually 80)
- The server user domain

For the MM7 content interface you require:

- The message type
- **submit.REQ** to send a notification message to the sender in the form of a submit request. The message goes from a VAS application to the MMSC.
- **deliver.REQ** to send a notification message to the sender in the form of a deliver request. The message goes from the MMSC to a VAS application.
- The hostname of the server
- The URL of the server (usually "/")
- The server port (usually 80)
- A user name and password to connect to the server
- The value-added-service-provider (VASP) ID
- The value-added-service (VAS) ID

For more information, see MMS notifications.

To configure administrator alert notifications - web-based manager

1. Go to **Firewall Objects > MMS Profile** and edit or add a new MMS protection profile.
2. Expand **MMS Bulk Email Filtering Detection**.
There are three message flood thresholds.
3. Expand the threshold that you want to configure alert notification for.
4. For **Message Flood Action**, select the **Alert Notification** check box. Alert notification options appear.
5. For the **Source MSISDN**, enter the MSISDN from which the alert notification message will be sent.
6. Select the Message Protocol the alert notification will use: **MM1**, **MM3**, **MM4**, or **MM7**.
7. Add the information required by FortiOS Carrier to send messages using the selected message protocol:
8. For **Notifications Per Second Limit**, enter the number of notifications to send per second.
Use this setting to reduce control the number of notifications sent by the FortiOS Carrier unit. If you enter zero (0), the notification rate is not limited.
9. If required, change **Window Start Time** and **Window Duration** configure when the FortiOS Carrier unit sends alert notifications.
By default, notifications are sent at any time of the day. You can change the Window Start Time if you want to delay sending alert messages. You can also reduce the Window Duration if you want to stop sending alert notifications earlier.

For example, you might not want FortiOS Carrier sending notifications except during business hours. In this case the Window Start Time could be 9:00 and the Window Duration could be 8:00 hours.

You can set different alert notifications for each message threshold. For example, you could limit the message window for lower thresholds and set it to 24 hours for higher thresholds. This way administrators will only receive alert notifications outside of business hours for higher thresholds.
10. For **Day of Week**, select the days of the week to send notifications.
For example, you may only want to send alert notifications on weekends for higher thresholds.
11. In the **Interval field**, enter the maximum frequency that alert notification messages will be sent, in minutes or hours.
All alerts occurring during the interval will be included in a single alert notification message to reduce the number of alert messages that are sent.

Configuring who to send alert notifications to

In each MMS protection profile you add a list of recipient MSISDNs. For each of these MSISDNs you select the message flood threshold that triggers sending notifications to this MSISDN.

To configure the alert notification recipients - web-based manager

1. Go to **Firewall Objects > MMS Profile**.
2. Select the **Edit** icon of the MMS profile in which you want to configure the alert notification recipients.
3. Expand **MMS Bulk Email Filtering Detection**.
4. Expand **Recipient MSISDN**.
5. Select **Create New**.
6. In the **New MSISDN** window, enter the MSISDN to use for flood threshold alert notification.
7. Select the duplicate thresholds at which to send alert notifications to the MSISDN.



For the flood threshold to be able to send an alert notification to the MSISDN, the alert notification action must be enabled and configured within the flood threshold.


Duplicate message protection

The convenience offered by MM1 and MM4 messaging can be abused by users sending spam or other unwanted messages. Often, the same message will be sent by multiple subscribers. The message can be spam, viral marketing, or worm-generated messages. MMS duplicate prevention can help prevent this type of abuse by keeping track of the messages being sent.

Overview

Duplicate message protection for MM1 messages prevents multiple subscribers from sending duplicate messages to your MMSC. Duplicate message protection for MM4 messages prevents another service provider from sending duplicate messages from the same subscriber to your MMSC. This can help prevent a potential flood that would otherwise become widespread between carriers.

MM1 and MM4 duplicate message protection



The FortiOS Carrier unit keeps track of the sent messages. If the same message appears more often than the threshold value you configure, then action is taken. Possible actions are logging the duplicates, blocking or intercepting duplicate messages, archiving the duplicate messages, and sending an alert to inform an administrator that duplicates are occurring.

With this highly configurable system, you can prevent the transmission of duplicate messages when there are more than you determine is acceptable.

For detailed configuration options, see [Duplicate Message](#).

Using message fingerprints to identify duplicate messages

The Carrier-enabled FortiGate unit detects duplicates by keeping a record of all the messages travelling on the network and comparing new messages to those that have already been sent.

Rather than save the messages, the FortiOS carrier creates a checksum using the message body and subject. This serves as a fingerprint to identify the message. If another message with the same message body and subject appears, the fingerprint will also be the same and the Carrier-enabled FortiGate unit will recognize it as a duplicate.

By creating and saving message fingerprints instead of saving the messages, the Carrier-enabled FortiGate unit can save resources and time.

Messages from any sender to any recipient

Duplicate message detection will detect duplicate messages regardless of the sender or recipient. To do this, message fingerprints are generated using only the message body and subject. The sender, recipient, and other header information is not included.

If multiple messages appear with the same subject and message body, the Carrier-enabled FortiGate unit will recognize them as being the same.

Setting duplicate message thresholds

The FortiOS Carrier recognizes all duplicate messages, but it will take action when it detects a volume of duplicate messages that exceed the duplicate threshold you set. The threshold defines the maximum number of duplicate messages allowed, the period during which the messages are considered, and the length of time the duplicate message can not be sent by anyone.

For example, you may determine that once a duplicate message is sent more than 300 times in an hour, any attempt to send the same duplicate message will be blocked for 30 minutes.

If a particular duplicate message exceeds the duplicate message threshold and is blocked, any further attempts to send the same message will re-start the block period.

Using the example above, if the duplicate message count exceeds the duplicate threshold, any attempt to send a copy of the duplicate message will be blocked for 30 minutes. If a subscriber tries to send a copy of the message after waiting 15 minutes, the message will be blocked and the block period will be reset to 30 minutes. The block period must expire with no attempts to send a duplicate message. Only then will a subscriber be allowed to send the message. Non-duplicate messages will not reset the block period.

Duplicate message actions

When the Carrier-enabled FortiGate unit detects that a duplicate message has exceeded duplicate threshold, it can take any combination of the five actions you configure for the duplicate threshold.

Action	Description
Log	Add a log entry indicating that a duplicate message event has occurred. You must also enable logging for MMS Scanning > Bulk Messages in the Logging section of the MMS protection profile.
DLP Archive	

Action	Description
All messages	Save all the messages that exceed the duplicate threshold in the DLP archive.
First message only	Save the first message to exceed the duplicate threshold in the DLP archive. Subsequent messages that exceed the duplicate threshold will not be saved.
Intercept	Messages that exceed the duplicate threshold are passed to the recipients, but if quarantine is enabled for intercepted messages, a copy of each message is also quarantined for later examination. If the quarantine of intercepted messages is disabled, the Intercept action has no effect.
Block	Messages that exceed the duplicate threshold are blocked and will not be delivered to the message recipients. If quarantine is enabled for blocked messages, a copy of each blocked message is quarantined for later examination.
Alert Notification	If the duplicate threshold is exceeded, the Carrier-enabled FortiGate unit will send an MMS duplicate message notification message.

Notifying duplicate message senders and receivers

The FortiOS Carrier unit does not send notifications to the sender or receiver of duplicate messages. If the sender or receiver is an attacker and is explicitly informed that they have exceeded a message threshold, the attacker may try to determine the exact threshold value by trial and error and then find a way around duplicate message protection. For this reason, no notification is set to the sender or receiver.

However, the FortiOS Carrier unit does have replacement messages for sending reply confirmations to MM1 senders and receivers and for MM4 senders for blocked messages identified as duplicate messages. For information about how FortiOS Carrier responds when message flood detection blocks a message, see and MMS duplicate messages and message floods.

Responses to MM1 senders and receivers

When the FortiOS Carrier unit identifies an MM1 message sent by a sender to an MMSC as a duplicate message and blocks it, the FortiOS Carrier unit returns a message submission confirmation (m-send.conf) to the sender (otherwise the sender's handset would keep retrying the message). The m-send.conf message is sent only when the MM1 duplicate message action is set to Block. For other duplicate message actions the message is actually delivered to the MMSC and the MMSC sends the m-send.conf message.

You can customize the m-send.conf message by editing the **MM1 send-conf duplicate message** MM1 replacement message (from the CLI the `mm1-send-conf-dupe` replacement message). You can customize the response status and message text for this message. The default response status is "Content not accepted". To hide the fact that the FortiOS Carrier unit is responding to a duplicate message, you can change the response status to "Success". The default message text informs the sender that the message was blocked. You could change this to something more generic.

For example, the following command sets the submission confirmation response status to "Success" and changes the message text to "Message Sent OK":

```
config system replacemsg mm1 mm1-send-conf-dupe
    set rsp-status ok
    set rsp-text "Message Sent OK"
end
```

When the FortiOS Carrier unit identifies an MM1 message received by a receiver from an MMSC as a duplicate message and blocks it, the FortiOS Carrier unit returns a message retrieval confirmation (m-retrieve.conf) to the sender (otherwise the sender's handset would keep retrying). The m-retrieve.conf message is sent only when the MM1duplicate message action is set to Block. For other message flood actions the message is actually received by the receiver, so the MMSC sends the m-retrieve.conf message.

You can customize the m-retrieve.conf message by editing the **MM1 retrieve-conf duplicate message** MM1 replacement message (from the CLI the mm1-retr-conf-dupe replacement message). You can customize the class, subject, and message text for this message.

For example, you could use the following command make the response more generic:

```
config system replacemsg mm1 mm1-retr-conf-dupe
    set subject "Message blocked"
    set message "Message temporarily blocked by carrier"
end
```

Forward responses for duplicate MM4 messages

When the FortiOS Carrier unit identifies an MM4 message as a duplicate message and blocks it, the FortiOS Carrier unit returns a message forward response (MM4_forward.res) to the forwarding MMSC (otherwise the forwarding MMSC would keep retrying the message). The MM4_forward.res message is sent only when the MM4 duplicate message action is set to Block and the MM4-forward.req message requested a response. For more information, see and MMS duplicate messages and message floods.

You can customize the MM4_forward.res message by editing the **MM4 duplicate message** MM4 replacement message (from the CLI the mm4-dupe replacement message). You can customize the response status and message text for this message. The default response status is "Content not accepted" (err-content-not-accept). To hide the fact that the FortiOS Carrier unit is responding to a duplicate message, you can change the response status to "Success". The default message text informs the sender that the message was blocked. You could change this to something more generic.

For example, the following command sets the submission confirmation response status to "Success" and changes the message text to "Message Forwarded OK":

```
config system replacemsg mm4 mm4-dupe
    set rsp-status ok
    set rsp-text "Message Forwarded OK"
end
```

Viewing DLP archived messages

If **DLP Archive** is a selected duplicate message action, the messages that exceed the threshold are saved to the MMS DLP archive. The default behavior is to save all of the offending messages but you can configure the DLP archive setting to save only the first message that exceeds the threshold. See Viewing DLP archived messages.

Order of operations: flood checking before duplicate checking

Although duplicate checking involves only examination and comparison of message contents and not the sender or recipient, and flood checking involves only totalling the number of messages sent by each subscriber

regardless of the message content, there are times when a selection of messages exceed both flood and duplicate thresholds.

The Carrier-enabled FortiGate unit checks for message floods before checking for duplicate messages. Flood checking is less resource-intensive and if the flood threshold invokes a **Block** action, the blocked messages are stopped before duplicate checking occurs. This saves both time and FortiOS Carrier system resources.

Bypassing duplicate message detection based on user's carrier endpoints

You can use carrier endpoint filtering to exempt MMS sessions from duplicate message detection. Carrier endpoint filtering matches carrier endpoints in MMS sessions with carrier endpoint patterns. If you add a carrier endpoint pattern to a filter list and set the action to exempt from mass MMS, all messages from matching carrier endpoints bypass duplicate message detection. For more information about endpoints, see FortiOS Handbook User Authentication guide.

Configuring duplicate message detection

To have the Carrier-enabled FortiGate unit check for duplicate messages, configure the duplicate threshold in an MMS profile, and select the MMS profile in a security policy.

All traffic matching the security policy will be checked for duplicate messages according to the settings in the MMS profile.



The duplicate scanner will only scan content. It will not scan headers. Content must be exactly the same. If there is any difference at all in the content, it will not be considered a duplicate.

The modular nature of the profiles allows you great flexibility in how you configure the scanning options. MMS profiles can be used in any number of policies, with different GTP profiles.

In a complex configuration, there may be many security policies, each with a different MMS profile. For a simpler network, you may have many security policies all using the same MMS profile.

Sending administrator alert notifications

When duplicate messages are detected, the Carrier-enabled FortiGate unit can be configured to notify you immediately with an MMS message. Enable this feature by selecting Alert Notification in the duplicate message action. Each duplicate message threshold can be configured separately.

Configuring how and when to send alert notifications

You can configure different alert notifications for MM1 and MM4 duplicate messages. You can configure the FortiOS Carrier unit to send these alert notifications using the MM1, MM3, MM4, or MM7 content interface. Each of these content interfaces requires alert notification settings that the FortiOS Carrier unit uses to communicate with a server using the selected content interface.

For the MM1 content interface you require:

- The hostname of the server
- The URL of the server (usually "/")
- The server port (usually 80)

For the MM3 and MM4 content interfaces you require:

- The hostname of the server
- The server port (usually 80)
- The server user domain

For the MM7 content interface you require:

- The message type
- **submit.REQ** to send a notification message to the sender in the form of a submit request. The message goes from a VAS application to the MMSC.
- **deliver.REQ** to send a notification message to the sender in the form of a deliver request. The message goes from the MMSC to a VAS application.
- The hostname of the server
- The URL of the server (usually "/")
- The server port (usually 80)
- A user name and password to connect to the server
- The value-added-service-provider (VASP) ID
- The value-added-service (VAS) ID

To configure administrator alert notifications - web-based manager

1. Go to **Security Profiles > MMS Profile** and edit or add a new MMS protection profile.
2. Expand **MMS Bulk Email Filtering Detection**.
There are three duplicate message thresholds.
3. Expand the threshold that you want to configure alert notification for.
4. For **Duplicate Message Action**, select the **Alert Notification** check box. Alert notification options appear.
5. For the **Source MSISDN**, enter the MSISDN from which the alert notification message will be sent.
6. Select the Message Protocol the alert notification will use: **MM1**, **MM3**, **MM4**, or **MM7**.
7. Add the information required by FortiOS Carrier to send messages using the selected message protocol:
8. For **Notifications Per Second Limit**, enter the number of notifications to send per second.
Use this setting to reduce control the number of notifications sent by the FortiOS Carrier unit. If you enter zero (0), the notification rate is not limited.
9. If required, change **Window Start Time** and **Window Duration** configure when the FortiOS Carrier unit sends alert notifications.
By default, notifications are sent at any time of the day. You can change the Window Start Time if you want to delay sending alert messages. You can also reduce the Window Duration if you want to stop sending alert notifications earlier.

For example, you might not want FortiOS Carrier sending notifications except during business hours. In this case the Window Start Time could be 9:00 and the Window Duration could be 8:00 hours.

You can set different alert notifications for each message threshold. For example, you could limit the message window for lower thresholds and set it to 24 hours for higher thresholds. This way administrators will only receive alert notifications outside of business hours for higher thresholds.
10. For **Day of Week**, select the days of the week to send notifications.
For example, you may only want to send alert notifications on weekends for higher thresholds.
11. In the **Interval field**, enter the maximum frequency that alert notification messages will be sent, in minutes or hours.

All alerts occurring during the interval will be included in a single alert notification message to reduce the number of alert messages that are sent.

Configuring who to send alert notifications to

In each MMS protection profile you add a list of recipient MSISDNs. For each of these MSISDNs you select the duplicate threshold that triggers sending notifications to this MSISDN.

To configure the alert notification recipients - web-based manager

1. Go to **Security Profiles > MMS Profile**.
2. Select the **Edit** icon of the MMS profile in which you want to configure the alert notification recipients.
3. Expand **MMS Bulk Email Filtering Detection**.
4. Expand **Recipient MSISDN**.
5. Select **Create New**.
6. In the **New MSISDN** window, enter the MSISDN to use for duplicate threshold alert notification.

Select the duplicate thresholds at which to send alert notifications to the MSISDN.



For the duplicate threshold to be able to send an alert notification to the MSISDN, the duplicate message threshold alert notification action must be enabled and configured.

Employing MMS Security features

FortiOS Carrier includes all the Security features of FortiOS with extra features specific to MMS carrier networks.

This section includes:

Why scan MMS messages for viruses and malware?

The requirement for scanning MM1 content comes from the fact that MMS is an increasingly popular technique for propagating malware between mobile devices.

Example: COMMWARRIOR

This is a virus for Series 60 type cell phones, such as Nokia, operating Symbian OS version 6 [or higher]. The object of the virus is to spread to other phones using Bluetooth and MMS as transport avenues. The targets are selected from the contact list of the infected phone and also sought via Bluetooth searching for other Bluetooth-enabled devices (phones, printers, gaming devices etc.) in the proximity of the infected phone.

This virus is more than a proof of concept - it has proven successfully its ability to migrate from a zoo collection to being in-the-wild. Currently, this virus is being reported in over 18 different countries around Europe, Asia and North America.

When the virus first infects a cell phone, a prompt is displayed asking the recipient if they want to install "Caribe". Symptoms of an infected phone may include rapid battery power loss due to constant efforts by the virus to spread to other phones via a Bluetooth seek-and-connect outreach.

The following variants among others are currently scanned by the FortiOS Carrier devices, in addition to more signatures that cover all known threats.

- **SymbOS/COMWAR.V10B!WORM**

- Aliases: SymbOS.Commwarrior.B, SymbOS/Commwar.B, SymbOS/Commwar.B!wm, SymbOS/Commwar.B-net, SymbOS/Commwarrior.b!sis, SymbOS/Comwar.B, SymbOS/Comwar.B!wm, SymbOS/Comwar.B-wm, SYMBOS_COMWAR.B, SymbOS/Comwar.1.0.B!wormSYMBOS/COMWAR.V10B.SP!WORM [Spanish version]
- First Discovered In The Wild: July 04, 2007
- Impact Level: 1
- Virus Class: Worm
- Virus Name Size: 23,320

- **SymbOS/Commwar.A!worm**

- Aliases: Commwarrior-A, SymbOS.Commwarrior.A [NAV], SymbOS/Commwar.A-net, SymbOS/Commwar_ezboot.A-ne, SymbOS/Comwar.A, SymbOS/Comwar.A-wm, SYMBOS_COMWAR.A [Trend]
- First Discovered In The Wild: May 16 2005
- Impact Level: 1
- Virus Class: Worm
- Virus Name Size: 27,936
- SymbOS/Commwarriie.C-wm
- Aliases: None
- First Discovered In The Wild: Oct 17 2005
- Impact Level: 1
- Virus Class: File Virus
- Virus Name Size: None

For the latest list of threats Fortinet devices detect, visit the FortiGuard Center.

MMS virus scanning

You can use MMS virus scanning to scan content contained within MMS messages for viruses. FortiOS Carrier virus scanning can be applied to the MM1, MM3, MM4, and MM7 interfaces to detect and remove content containing viruses at many points in an MMS network. Perhaps the most useful interface to apply virus scanning would be the MM1 interface to block viruses sent by mobile users before they get into the service provider network.

To go to MMS virus scanning, go to **Security Profiles MMS Profile**, select an existing or create a new profile, and expand **MMS Scanning**. See MMS scanning options.

This section includes:

- [MMS virus monitoring](#)
- [MMS virus scanning blocks messages \(not just attachments\)](#)
- [Scanning MM1 retrieval messages](#)
- [Configuring MMS virus scanning](#)
- [Removing or replacing blocked messages](#)
- [Carrier Endpoint Block](#)
- [MMS Content Checksum](#)
- [Passing or blocking fragmented messages](#)
- [Client comforting](#)
- [Server comforting](#)
- [Handling oversized MMS messages](#)

MMS virus monitoring

To enable MMS virus monitoring, expand **MMS Scanning** and enable **Monitor only** for the selected MMS types.

This feature causes the FortiOS Carrier unit to record log messages when MMS scanning options find a virus, match a file name, or match content using any of the other MMS scanning options. Selecting this option enables reporting on viruses and other problems in MMS traffic without affecting users.

MMS virus scanning blocks messages (not just attachments)

To enable MMS virus scanning, expand **MMS Scanning** and enable **Virus Scan** for the selected MMS types.

Because MM1 and MM7 use HTTP, the oversize limits for HTTP and the HTTP antivirus port configurations also apply to MM1 and MM7 scanning. See

MM3 and MM4 use SMTP and the oversize limits for SMTP and the SMTP antivirus port configurations also apply to MM3 and MM4 scanning.

The message contents will be scanned for viruses, matched against the file extension blocking lists and scanned for banned words. All these items will be configured via the standard GUI interfaces available for the other protocols and will be controlled at the protection profile level with new options specifically for the MM1 messages.

The FortiOS Carrier unit extracts the sender's Mobile Subscriber Integrated Services Digital Network Number (MSISDN) from the HTTP headers if available. The `POST` payload will be sent to the scan units which will parse the MMS content and scan each message data section. If any part of the data is to be blocked, the proxy will be informed, the connection to the MMSC will be reset and the Carrier-enabled FortiGate unit will return an `HTTP 200 OK` message with an `m-send-conf` payload to the client to prevent a retry. Finally the appropriate logging, alert, and replacement message events will be triggered.

For client notification, the `x-mms-response-status` and `x-mms-response-text` fields can also be customized as required.

Scanning MM1 retrieval messages

To scan MM1 retrieval messages, expand **MMS Scanning** and select **Scan MM1 message retrieval**.

Select to scan message retrievals that use MM1. If you enable **Virus Scan** for all MMS interfaces, messages are also scanned while being sent. In this case, you can disable MM1 message retrieval scanning to improve performance.

Configuring MMS virus scanning

To configure MMS virus scanning, expand **MMS Scanning** and enable **Virus Scan**.

Once applied to a security policy, the MMS protection profile will then perform virus scans on all traffic accepted by that policy.

Removing or replacing blocked messages

To remove blocked messages, expand **MMS Scanning** and select **Remove Blocked** for the selected MMS types.

Select **Remove Blocked** remove blocked content from each protocol and replace it with the replacement message. If FortiOS Carrier is to preserve the length of the message when removing blocked content, as may occur when billing is affected by the length of the message, select **Constant**.

If you only want to monitor blocked content, select **Monitor Only**.

Carrier Endpoint Block

A carrier endpoint defines a specific client on the carrier network. Typically the client IP address is used to identify the client, however on a carrier network this may be impractical when the client is using a mobile device. Other identifying information such as the MSISDN number is used instead.

This information can be used to block a specific endpoint on the network. Reasons for blocking may include clients whose accounts have expired, clients from another carrier, clients who have sent malicious content (phishing, exploits, viruses, etc), or other violations of terms of use.

Enabling carrier endpoint blocking

To enable carrier endpoint blocking you first need to create a carrier endpoint filter list, and then enable it.

To enable carrier endpoint blocking - web-based manager

1. Create a carrier endpoint filter list.
2. Go to **Security Profiles > MMS Profile**.
3. Select **Create New**, or select an existing profile to edit and select **Edit**.
4. Expand MMS Scanning.
5. Select one or more types of MMS messaging to enable endpoint blocking on.
6. Select the carrier endpoint filter list to use in matching the endpoints to be blocked.



In MMS Profile, endpoints can only be blocked.

Create a carrier endpoint filter list

A carrier endpoint filter list contains one or more carrier endpoints to match. When used in MMS scanning entries in the filter list that are matched are blocked.

You can configure multiple filter lists for different purposes and groups of clients, such as blocking clients, clients with different levels of service agreements, and clients from other carriers. See Carrier endpoint filter lists configuration settings.

To create a carrier endpoint filter list - web-based manager

1. Go to **Security Profiles > Carrier Endpoint Filter Lists**.
2. Select **Create New**.
3. Enter a descriptive name for the filter list, such as `blocked_clients` or `CountryX_clients`, and select **OK**.
4. Select **Create New** to add one or more entries to the list.
5. Select **OK** to return to display the list of filter lists.

Configuring endpoint filter list entries

For each single endpoint or group of endpoints have part of their identifying information in common, you create an entry in the endpoint filter list.

For example a `blocked_clients` filter list may include entries for single endpoints added as each one needs to be blocked and a group of clients from a country that does not allow certain services.

To configure an endpoint filter list entry - web-based manager

1. Select **Create New**.
2. Enter the following information and select **OK**.

Name	Name of endpoint filter list. Select this name in an MMS protection profile.
Comments	Optional description of the endpoint filter list.
Check/Uncheck All	<p>Select the check box to enable all endpoint patterns in the MMS filter list.</p> <p>Clear the check box to disable all entries on the MMS filter list.</p> <p>You can also select or clear individual check boxes to enable or disable individual endpoint patterns.</p>
Pattern	The pattern that FortiOS Carrier uses to match with endpoints. The pattern can be a single endpoint or consist of wildcards or Perl regular expressions that will match more than one endpoint. For more on wildcard and regular expressions, see Using wildcards and Perl regular expressions in the UTM guide.
Action	<p>Select the action taken by FortiOS Carrier for messages from a carrier endpoint that matches the endpoint pattern:</p> <p>None - No action is taken.</p> <p>Block - MMS messages from the endpoint are not delivered and FortiOS Carrier records a log message.</p> <p>Exempt from mass MMS - MMS messages from the endpoint are delivered and are exempt from mass MMS filtering. Mass MMS filtering is configured in MMS protection profiles and is also called MMS Bulk Email Filtering and includes MMS message flood protection and MMS duplicate message detection. A valid use of mass MMS would be when a service provider notifies customers of a system-wide event such as a shutdown.</p> <p>Exempt from all scanning - MMS messages from the endpoint are delivered and are exempt from all MMS protection profile scanning.</p>

Content Archive	MMS messages from the endpoint are delivered, the message content is DLP archived according to MMS DLP archive settings. Content archiving is also called DLP archiving.
Intercept	MMS messages from the endpoint are delivered. Based on the quarantine configuration, attached files may be removed and quarantined.
Pattern Type	The pattern type: Wildcard , Regular Expression , or Single Endpoint .
Enable	Select to enable this endpoint filter pattern.

Blocking network access based on endpoints

You can use endpoint IP filtering to block traffic from source IP addresses associated with endpoints. You can also configure FortiOS Carrier to record log messages whenever endpoint IP filtering blocks traffic. Endpoint IP filtering blocks traffic at the IP level, before the traffic is accepted by a security policy.

To configure endpoint IP filtering, go to **Security Profiles > IP Filter** and add endpoints to the IP filter list. For each endpoint you can enable or disable both blocking traffic and logging blocked traffic.



You cannot add endpoint patterns to the endpoint IP filter list. You must enter complete and specific endpoints that are valid for your network.



The only action available is block. You cannot use endpoint IP filtering to exempt endpoints from IP filtering or to content archive or quarantine communication sessions.

FortiOS Carrier looks in the current user context list for the endpoints in the IP filter list and extracts the source IP addresses for these endpoints. Then any communication session with a source IP address that matches one of these IP addresses is blocked at the IP level, before the communication session is accepted by a security policy.

FortiOS Carrier dynamically updates the list of IP addresses to block as the user context list changes. Only these updated IP addresses are blocked by endpoint IP filtering.

For information about the carrier endpoints and the user context list, including how entries are added to and removed from this list.

MMS Content Checksum

The MMS content checksum feature attempts to match checksums of known malicious MMS messages, and on a successful match it will be blocked. The checksums are applied to each part of the message—attached files and message body have separate checksums. These checksums are created with CRC-32, the same method as FortiAnalyzer checksums.

For example, if an MMS message contains a browser exploit in the message body, you can add the checksum for that message body to the list, and future occurrences of that exact message will be blocked. Content will be replaced by the content checksum block notification replacement message for that type of MMS message, and if it is enabled the event will be logged.

One possible implementation would be to configure all .sis files to be intercepted. When one is found to be infected or malicious it would be added to the MMS content checksum list.

To use this feature a list of one or more malicious checksums must be created and then the feature is enabled using that list. For a detailed list of options, see MMS Content Checksum.

To configure an MMS content checksum list

1. Go to **Security Profiles > MMS Content Checksum**.
2. Select **Create New**.
3. Enter a name for the list of checksums, and select **OK**.
You are taken to the edit screen for that new list.
4. Select **Create New** to add a checksum.
5. Enter the **Name** and **Checksum**, and select **OK**.
The checksum is added to the list.

To add more checksums to the list, repeat steps 4 and 5.

To remove a checksum from the list you can either delete the checksum or simply disable it and leave it in the list.

To enable MMS content checksums, expand **MMS Scanning** and select **MMS Content Checksum** for the selected MMS types. Select the checksum list to match.

Passing or blocking fragmented messages

Select to pass fragmented MM3 and MM4 messages. Fragmented MMS messages cannot be scanned for viruses. If you do not select these options, fragmented MM3 and MM4 messages are blocked.

The **Interval** is the time in seconds before client comforting starts after the download has begun, and the time between sending subsequent data.

The **Amount** is the number of bytes sent by client or server comforting at each interval.

Client comforting

In general, client comforting is available for MM1 and MM7 messaging and provides a visual display of progress for web page loading or HTTP or FTP file downloads. Client comforting does this by sending the first few packets of the file or web page being downloaded to the client at configured time intervals so that the client is not aware that the download has been delayed. The client is the web browser or FTP client. Without client comforting, clients and their users have no indication that the download has started until the Carrier-enabled FortiGate unit has completely buffered and scanned the download. During this delay users may cancel or repeatedly retry the transfer, thinking it has failed.

The appearance of a client comforting message (for example, a progress bar) is client-dependent. In some instances, there will be no visual client comforting cue.

During client comforting, if the file being downloaded is found to be infected, then the Carrier-enabled FortiGate unit caches the URL and drops the connection. The client does not receive any notification of what happened because the download to the client had already started. Instead the download stops, and the user is left with a partially downloaded file.

If the user tries to download the same file again within a short period of time, then the cached URL is matched and the download is blocked. The client receives the Infection cache message replacement message as a

notification that the download has been blocked. The number of URLs in the cache is limited by the size of the cache.



Client comforting can send unscanned (and therefore potentially infected) content to the client. Only enable client comforting if you are prepared to accept this risk. Keeping the client comforting interval high and the amount low will reduce the amount of potentially infected data that is downloaded.

MM1 and MM7 client comforting steps

Since MM1 and MM7 messages use HTTP, MM1 and MM7 client comforting operates like HTTP client comforting.

The following steps show how client comforting works for a download of a 1 Mbyte file with the client comforting interval set to 20 seconds and the client comforting amount set to 512 bytes.

1. The client requests the file.
2. The Carrier-enabled FortiGate unit buffers the file from the server. The connection is slow, so after 20 seconds about one half of the file has been buffered.
3. The Carrier-enabled FortiGate unit continues buffering the file from the server, and also sends 512 bytes to the client.
4. After 20 more seconds, the FortiGate unit sends the next 512 bytes of the buffered file to the client.
5. When the file has been completely buffered, the client has received the following amount of data:

$$ca * (T/ci) \text{ bytes} == 512 * (40/20) == 512 * 2 == 1024 \text{ bytes,}$$
 where *ca* is the client comforting amount, *T* is the buffering time and *ci* is the client comforting interval.
6. If the file does not contain a virus, the Carrier-enabled FortiGate unit sends the rest of the file to the client. If the file is infected, the FortiGate closes the data connection but cannot send a message to the client.

Server comforting

Server comforting can be selected for each protocol.

Similar to client comforting, you can use server comforting to prevent server connection timeouts that can occur while waiting for FortiOS Carrier to buffer and scan large `POST` requests from slow clients.

The **Interval** is the time in seconds before client and server comforting starts after the download has begun, and the time between sending subsequent data.

The **Amount** is the number of bytes sent by client or server comforting at each interval.

Handling oversized MMS messages

Select **Block** or **Pass** for files and email messages exceeding configured thresholds for each protocol.

The oversize threshold refers to the final size of the message, including attachments, after encoding by the client. Clients can use a variety of encoding types; some result in larger file sizes than the original attachment. As a result, a file may be blocked or logged as oversized even if the attachment is several megabytes smaller than the oversize threshold.

MM1 sample messages

```
Internet Protocol, Src Addr: 10.128.206.202 (10.128.206.202), Dst Addr: 10.129.192.190
(10.129.192.190)
Transmission Control Protocol, Src Port: 34322 (34322), Dst Port: http (80), Seq: 1, Ack:
1, Len: 1380
Source port: 34322 (34322)
Destination port: http (80)
Header length: 20 bytes
Flags: 0x0010 (ACK)
Window size: 24840
Checksum: 0x63c1 (correct)
```

HTTP proxy

```
Hypertext Transfer Protocol
POST / HTTP/1.1\r\n
Request Method: POST
Request URI: /
Request Version: HTTP/1.1
Host: 10.129.192.190\r\n
Accept: /*, application/vnd.wap.sic,application/vnd.wap.mms-message,text/x-
html,image/mng,image/x-mng,video/mng,video/x-mng,image/bmp\r\n
Accept-Charset: utf-8,*\r\n
Accept-Language: en\r\n
Content-Length: 25902\r\n
Content-Type: application/vnd.wap.mms-message\r\n
User-Agent: Nokia7650/1.0 SymbianOS/6.1 Series60/0.9 Profile/MIDP-1.0
Configuration/CLDC-1.0 UP.Link/6.2.1\r\n
x-up-devcap-charset: utf-8\r\n
x-up-devcap-max-pdu: 102400\r\n
x-up-uplink: magh-ip.mi.vas.omnitel.it\r\n
x-wap-profile: "http://nds.nokia.com/uaprof/N7650r200.xml"\r\n
x-up-subno: 1046428312-826\r\n
x-up-calling-line-id: 393475171234\r\n
x-up-forwarded-for: 10.211.4.12\r\n
x-forwarded-for: 10.211.4.12\r\n
Via: 1.1 magh-ip.mi.vas.omnitel.it\r\n
\r\n
```

Scan engine

```
MMS Message Encapsulation, Type: m-send-req
X-Mms-Message-Type: m-send-req (0x80)
X-Mms-Transaction-ID: 1458481935
X-Mms-MMS-Version: 1.0
From: <insert address>
To: 3475171234/TYPE=PLMN
X-Mms-Message-Class: Personal (0x80)
X-Mms-Expiry: 21600.000000000 seconds
X-Mms-Priority: Normal (0x81)
X-Mms-Delivery-Report: No (0x81)
X-Mms-Read-Report: No (0x81)
Content-Type: application/vnd.wap.multipart.related; start=<1822989907>;
type=application/smil
Start: <1822989907>
Type: application/smil
```

```

Data (Post)
  Multipart body
    Part: 1, content-type: text/plain
      Content-Type: text/plain; charset=iso-10646-ucs-2; name=Ciao.txt
      Charset: iso-10646-ucs-2
      Name: Ciao.txt
    Headers
      Content-Location: Ciao.txt
      Line-based text data: text/plain
        \377\376C\000i\000a\000o\000
[Unreassembled Packet: MMSE]

```

Sender notifications and logging

In most cases you will notify the sender that they are causing problems on the network — either by sending malware content, flooding the network, or some other unwanted activity. The notification assumes the sender is unaware of their activity and will stop or correct it when notified.

However, senders who are notified may use this information to circumvent administration's precautions. For example if flood notification is set to 1000 messages per minute, a notified user may simply reduce their message to 990 messages per minute if this flood is intentional. For this reason, not all problems include sender notifications.

There are two methods of notifying senders:

- [MMS notifications](#)
- [Replacement messages](#)

And three details to consider for logging and notifying administrators:

- [Logging and reporting](#)
- [MMS logging options](#)
- [SNMP](#)

MMS notifications

MMS notifications enable you to customize notifications for many different situations and differently for all the supported MMS message protocols — MM1, MM3, MM4, and MM7.

MMS notification types include:

- Content Filter
- File Block
- Carrier Endpoint Block
- Flood
- Duplicate
- MMS Content Checksum
- Virus Scan

Day of Week, **Window start time** and **Window Duration** define what days and what time of day alert notifications will be sent. This allows you to control what alerts are sent on weekends. It also lets you control when to start sending notifications each day. This can be useful if system maintenance is performed at the same time each night — you might want to start alert notifications after maintenance has completed. Another reason to limit the time alert messages are sent could be to limit message traffic to business hours.

Notifications screen for FortiOS Carrier MMS Profile

Edit MMS Profile																													
▶ MMS Bulk Email Filtering Detection																													
▶ MMS Address Translation																													
▼ MMS Notifications																													
		Option																											
AntiVirus Notification List		-- Disabled --																											
		MM1				MM3				MM4				MM7															
Message Protocol		mm1				mm3				mm4				mm7															
Message Type														deliver.REQ															
Detect Server Details		<input checked="" type="checkbox"/>				<input checked="" type="checkbox"/>				<input checked="" type="checkbox"/>				<input checked="" type="checkbox"/>															
Hostname																													
URL		/												/															
Port		80				25				25				80															
Username																													
Password																													
VASP ID																													
VAS ID																													
▶ All Notification Types		<input type="checkbox"/>	24	hour(s)		<input type="checkbox"/>	24	hour(s)		<input type="checkbox"/>	24	hour(s)		<input type="checkbox"/>	24	hour(s)													
Notifications Per Second Limit		0				0				0				0															
Day of Week		Sun	Mon	Tue	Wed	Thu	Fri	Sat	Sun	Mon	Tue	Wed	Thu	Fri	Sat	Sun	Mon	Tue	Wed	Thu	Fri	Sat	Sun	Mon	Tue	Wed	Thu	Fri	Sat
		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	
Window Start Time		00 : 00				00 : 00				00 : 00				00 : 00															
Window Duration		24 : 00				24 : 00				24 : 00				24 : 00															
▶ DLP Archive																													
▶ Logging																													
<div>OK</div> <div>Cancel</div>																													

For MMS Notification options, see MMS Notifications.

Replacement messages

FortiGate units send replacement messages when messages or content is blocked, quarantined, or otherwise diverted from the receiver. In its place a message is sent to notify the receiver what happened.

With FortiOS Carrier MMS replacement messages, send and receive message types are supported separately and receive their own custom replacement messages. This allows the network to potentially notify both the sender and receiver of the problem.

For example the replacement message **MM1 send-req file block message** is sent to the device that sent one or more files that were banned. The default message that is sent is `This device has sent %%NUM_MSG%% messages containing banned files in the last %%DURATION%% hours. The two variables are replaced by the appropriate values.`

Replacement messages are not as detailed or specific as MMS notifications, but they are also not as complicated to configure. They are also useful when content has been removed from an MMS message that was still delivered.

Logging and reporting

With each virus infection, or file block, a syslog message is generated. The format of this syslog message is similar to:

```
2005-09-22 19:15:47 deviceid=FGT5001ABCDEF1234 logid=0211060ABC type=virus
  subtype=infected level=warning src=10.1.2.3 dst=10.2.3.4 srcintf=port1 dstintf=port2
  service=mm1 status=blocked from="<sending MSISDN>" to="<receiving MSISDN>"
```



```
file="eicar.com.txt" virus="EICAR_TEST_FILE" msg="The file eicar.com.txt is infected
with EICAR_TEST_FILE. ref
http://www.fortinet.com/VirusEncyclopedia/search/encyclopediaSearch.do?method=quickSea
rchDirectly&virusName=EICAR_TEST_FILE"
```

Note that the **from** and **to** fields are samples and not real values.

MMS logging options

You can enable logging in an MMS protection profile to write event log messages when the MMS protection profile options that you have enabled perform an action. For example, if you enable MMS antivirus protection, you could also use the MMS protection profile logging options to write an event log message every time a virus is detected.

To record these log messages you must first configure how the FortiOS Carrier unit stores log messages.

To configure MMS content archiving, go to **Security Profiles > MMS Profile**. Select **Create New** or select the **Edit** icon beside an existing profile. Expand **MMS Bulk AntiSpam Detection > Logging**. Complete the fields as described in the following table and select **OK**. For more a detailed list of options, see Logging.

SNMP

A simple SNMP trap will be generated to inform the operators' alerting system that a virus has been detected. This SNMP trap could contain the sending and receiving MSISDN, however the initial solution would reflect the current behavior, i.e. only the fact that a virus has been detected will be communicated.

MMS content-based Antispam protection

Expand **MMS Scanning** and select **Content Filter** in an MMS protection profile to create content filter black/white lists that block or allow MMS messages based on the content of the message.

Overview

A school computer lab may block age-inappropriate content. A place of business may block unproductive content. A public access internet cafe may block offensive and graphic content. Each installation has its own requirements for what content needs to be blocked, and in what language.

FortiOS Carrier provides the ability to create custom local dictionaries, black lists, and white lists in multiple languages enables you to protect your customers from malicious content around the world.

Configurable dictionary

You can create a dictionary of configurable terms and phrases using the CLI. The text of MMS messages will be searched for these terms and phrases. Add content filter lists that contain content that you want to match in MMS messages. For every match found, a score is added. If enough matches are found to set the total score above the configured threshold, the MMS message is blocked.

You can add words, phrases, wild cards and Perl regular expressions to create content patterns that match content in MMS messages. For more on wildcard and regular expressions, see Using wildcards and Perl regular expressions in the UTM guide.

For each pattern you can select **Block** or **Exempt**.

- Block adds an antispam black list pattern. A match with a block pattern blocks a message depending on the score of the pattern and the content filter threshold.

- Exempt adds an antispam white list pattern. A match with an exempt pattern allows the message to proceed through the FortiOS Carrier unit, even if other content patterns in the same content filter list would block it.

If a pattern contains a single word, the FortiOS Carrier unit searches for the word in MMS messages. If the pattern contains a phrase, the FortiOS Carrier unit searches for all of the words in the phrase. If the pattern contains a phrase in quotation marks, the FortiOS Carrier unit searches for the whole phrase.

You can create patterns with Simplified Chinese, Traditional Chinese, Cyrillic, French, Japanese, Korean, Spanish, Thai, or Western character sets.

Black listing

Black listing is the practice of banning entries on the list. For example if an IP address continuously sends viruses, it may be added to the black list. That means any computers that consult that list will not communicate with that IP address.

Sometimes computers or devices can be added to black lists for a temporary problem, such as a virus that is removed when notified. However, as a rule short of contacting the administrator in person to manually be removed from the black list, users have to wait and they generally will be removed after a period without problem.

White listing

White listing is the practice of adding all critical IP addresses to a list, such as company email and web servers. Then if those servers become infected and start sending spam or viruses, those servers are not blocked. This allows the critical traffic through, even if there might be some malicious traffic as well. Blocking all traffic from your company servers would halt company productivity.

Scores and thresholds

Each content pattern includes a score. When a MMS message is matched with a pattern the score is recorded. If a message matches more than one pattern or matches the same pattern more than once, the score for the message increases. When the total score for a message equals or exceeds the threshold the message is blocked.

The default score for a content filter list entry is 10 and the default threshold is 10. This means that by default a message is blocked by a single match. You can change the scores and threshold so that messages can only be blocked if there are multiple matches. For example, you may only want to block messages that contain the phrase "example" if it appears twice. To do this, add the "example" pattern, set action to block and score to 5. Keep the threshold at 10. If "example" is found twice or more in a message the score adds up 10 (or more) and the message is blocked.

Configuring content-based antispam protection

To apply content-based antispam protection - CLI

```
config webfilter content
  edit <filter_table_number>
    set name <filter_table_name>
    config entries
      edit <phrase or regexp you want to block>
        set action {block | exempt}
        set lang <phrase language>
        set pattern-type {wildcard | regexp}
        set score <phrase score>
        set status {enable | disable}
      end
    end
```

end

Configuring sender notifications

When someone on the MMS network sends an MMS message that is blocked, in most cases you will notify the sender. Typically an administrator is notified in addition to the sender so action can be taken if required. There are two types of sender notifications available in FortiOS Carrier: MMS notifications, and Replacement Messages.

MMS notifications

MMS notifications to senders are configured in **Security Profiles > MMS Profile**, under MMS Notifications.

In this section you can configure up to four different notification recipients for any combination of MM1/3/4/7 protocol MMS messages. Also for MM7 messages the message type can be `submit.REQ` or `deliver.REQ`.

Useful settings include:

- delay in message based on notification type
- limit on notifications per second to prevent a flood
- schedules for notifications
- log in details for MM7 messages.

For more information on MMS notifications, see Notifying message flood senders and receivers and MMS Notifications.

Replacement messages

Replacement messages are features common to both FortiOS and FortiOS Carrier, however FortiOS Carrier has additional messages for the MMS traffic.

While each MMS protocol has its own different replacement messages, the one common to all MMS protocols is the **MMS blocked content replacement message**. This is the message that the receiver of the message sees when their content is blocked.

MMS DLP archiving

You can use DLP archiving to collect and view historical logs that have been archived to a FortiAnalyzer unit or the FortiGuard Analysis and Management service. DLP archiving is available for FortiAnalyzer when you add a FortiAnalyzer unit to the FortiOS Carrier configuration. The FortiGuard Analysis and Management server becomes available when you subscribe to the FortiGuard Analysis and Management Service.

You can configure full DLP archiving and summary DLP archiving. Full DLP archiving includes all content, for example, full email DLP archiving includes complete email messages and attachments. Summary DLP archiving includes just the meta data about the content, for example, email message summary records include only the email header.

You can archive MM1, MM3, MM4, and MM7 content.

Configuring MMS DLP archiving

Select DLP archive options to archive MM1, MM3, MM4, and MM7 sessions. For each protocol you can archive just session metadata (**Summary**), or metadata and a copy of the associated file or message (**Full**).

In addition to MMS protection profile DLP archive options you can:

- Archive MM1 and MM7 message floods
- Archive MM1 and MM7 duplicate messages
- Select **DLP archiving** for carrier endpoint patterns in a **Carrier Endpoint List** and select the **Carrier Endpoint Block** option in the **MMS Scanning** section of an MMS Protection Profile

FortiOS Carrier only allows one sixteenth of its memory for transferring content archive files. For example, for Carrier-enabled FortiGate units with 128 MB RAM, only 8 MB of memory is used when transferring content archive files. Best practices dictate to not enable full content archiving if antivirus scanning is also configured because of these memory constraints.

To configure MMS DLP archiving - web-based manager

1. Go to **Security Profiles > MMS Profile**.
2. Select **Create New** or select the **Edit** icon beside an existing profile.
3. Expand **MMS Bulk AntiSpam Detection > Content Archive**.
4. Complete the fields as described in DLP Archive options.
5. Select **OK**.

Viewing DLP archives

You can view DLP archives from the Carrier-enabled FortiGate unit web-based manager. Archives are historical logs that are stored on a log device that supports archiving, such as a FortiAnalyzer unit.

These logs are accessed from either **Log & Report > DLP Archive** or if you subscribed to the FortiCloud service, you can view log archives from there.

The **DLP Archive** menu is only visible if one of the following is true.

- You have configured the FortiGate unit for remote logging and archiving to a FortiAnalyzer unit.
- You have subscribed to FortiCloud.

The following tabs are available when you are viewing DLP archives for one of these protocols.

- **E-mail** to view POP3, IMAP, SMTP, POP3S, IMAPS, SMTPS, and spam email archives.
- **Web** to view HTTP and HTTPS archives.
- **FTP** to view FTP archives.
- **IM** to view AIM, ICQ, MSN, and Yahoo! archives.
- **MMS** to view MMS archives.
- **VoIP** to view session control (SIP, SIMPLE and SCCP) archives.

If you need to view log archives in Raw format, select **Raw** beside the **Column Settings** icon.

GTP basic concepts

GPRS currently supports data rates from 9.6 kbps to more than 100 kbps, and is best suited for burst forms of traffic. GPRS involves both radio and wired components. The mobile phone sends the message to a base station unit (radio based), and the base station unit sends the message to the carrier network and eventually the Internet (wired carrier network).

The network system then either sends the message back to a base station and to the destination mobile unit, or forwards the message to the proper carrier's network where it gets routed to the mobile unit.

PDP Context

The packet data protocol (PDP) context is a connection between a mobile station and the end address that goes through the SGSN and GGSN. It includes identifying information about the mobile customer used by each server or device to properly forward the call data to the next hop in the carrier network, typically using a GTP tunnel between the SGSN and GGSN.

When a mobile customer has an active voice or data connection open, both the SGSN and GGSN have the PDP context information for that customer and session.

When a mobile phone attempts to communicate with an address on an external packet network, either an IP or X.25 address, the mobile station that phone is connected to opens a PDP context through the SGSN and GGSN to the end address. Before any traffic is sent, the PDP context must first be activated.

The information included in the PDP context includes the customer's IP address, the IMSI number of the mobile handset, and the tunnel endpoint ID for both the SGSN and GGSN. The ID is a unique number, much like a session ID on a TCP/IP firewall. All this information ensures a uniquely identifiable connection is made.

Since one mobile device may have multiple connections open at one time, such as data connections to different Internet services and voice connections to different locations, there may be more than one PDP context with the same IP address making the extra identifying information required.

The endpoint that the mobile phone is connecting to only knows about the GGSN — the rest of the GPRS connection is masked by the GGSN.

Along the PDP context path, communication is accomplished in using three different protocols.

- The connection between the Mobile Station and SGSN uses the SM protocol.
- Between SGSN and GGSN GTP is used.
- Between GGSN and the endpoint either IP or X.25 is used.

FortiOS Carrier is concerned with the SGSN to GGSN part of the PDP context — the part that uses GTP.

For more about PDP context, see Tunnel Management Messages.

Creating a PDP context

While FortiOS Carrier is concerned mostly with the SGSN to GGSN part of the PDP Context, knowing the steps involved in creating a PDP context helps understand the role each device, protocol, and message type plays.

Both mobile stations and GGSNs can create PDP contexts.

A Mobile Station creates a PDP context

1. The Mobile Station (MS) sends a `PDP activation request` message to the SGSN including the MS PDP address, and APN.
2. Optionally, security functions may be performed to authenticate the MS.
3. The SGSN determines the GGSN address by using the APN identifier.
4. The SGSN creates a down link GTP tunnel to send IP packets between the GGSN and SGSN.
5. The GGSN creates an entry in its PDP context table to deliver IP packets between the SGSN and the external packet switching network.
6. The GGSN creates an uplink GTP tunnel to route IP-PDU from SGSN to GGSN.
7. The GGSN then sends back to the SGSN the result of the PDP context creation and if necessary the MS PDP address.

8. The SGSN sends an `Activate PDP context accept` message to the MS by returning negotiated the PDP context information and if necessary the MS PDP address.
9. Now traffic can pass from the MS to the external network endpoint.

A GGSN creates a PDP context

1. The network receives an IP packet from an external network.
2. The GGSN checks if the PDP Context has already been created.
3. If not, the GGSN sends a `PDU notification request` to the SGSN in order to initiate a PDP context activation.
4. The GGSN retrieves the IP address of the appropriate SGSN address by interrogating the HLR from the IMSI identifier of the MS.
5. The SGSN sends to the MS a request to activate the indicated PDP context.
6. The PDP context activation procedure follows the one initiated by the MS. See [“A Mobile Station creates a PDP context”](#).
7. When the PDP context is activated, the IP packet can be sent from the GGSN to the MS.

Terminating a PDP context

A PDP context remains open until it is terminated. To terminate the PDP context an MS sends a `Deactivate PDP context` message to the SGSN, which then sends a `Delete PDP Context` message to the GGSN. When the SGSN receives a PDP context deletion acknowledgment from the GGSN, the SGSN confirms to the MS the PDP context deactivation. The PDP can be terminated by the SGSN or GGSN as well with a slight variation of the order of the messages passed.

When the PDP Context is terminated, the tunnel it was using is deleted as well. If this is not completed in a timely manner, it is possible for someone else to start using the tunnel before it is deleted. This hijacking will result in the original customer being over billed for the extra usage. Anti-overbilling helps prevent this. See [Configuring Anti-overbilling in FortiOS Carrier](#).

GPRS security

The GPRS network has some built-in security in the form of GPRS authentication. However this is minimal, and is not sufficient for carrier network security needs. A GTP firewall, such as FortiOS Carrier, is required to secure the Gi, Gn, and Gp interfaces.

GPRS authentication

GPRS authentication is handled by the SGSN to prevent unauthorized GPRS calls from reaching the GSM network beyond the SGSN (the base station system, and mobile station). Authentication is accomplished using some of the customer's information with a random number and uses two algorithms to create ciphers that then allow authentication for that customer.

User identity confidentiality ensures that customer information stays between the mobile station and the SGSN — no identifying information goes past the SGSN. Past that point other numbers are used to identify the customer and their connection on the network.

Periodically the SGSN may request identity information from the mobile station to compare to what is on record, using the IMEI number.

Call confidentiality is achieved through the use of a cipher, similar to the GPRS authentication described earlier. The cipher is applied between the mobile station and the SGSN. Essentially a cipher mask is XORd with each outgoing frame, and the receiving side XORs with its own cipher to result in the original frame and data.

Parts of a GTPv1 network

A sample GTP network consists of the end handset sender, the sender's mobile station, the carrier's network including the SGSN and GGSN, the receiver's mobile station, and the receiver handset.

When a handset moves from one mobile station and SGSN to another, the handset's connection to the Internet using GTP tracks the user's location and information. For example, the handset could move from one cell to another, or between countries.

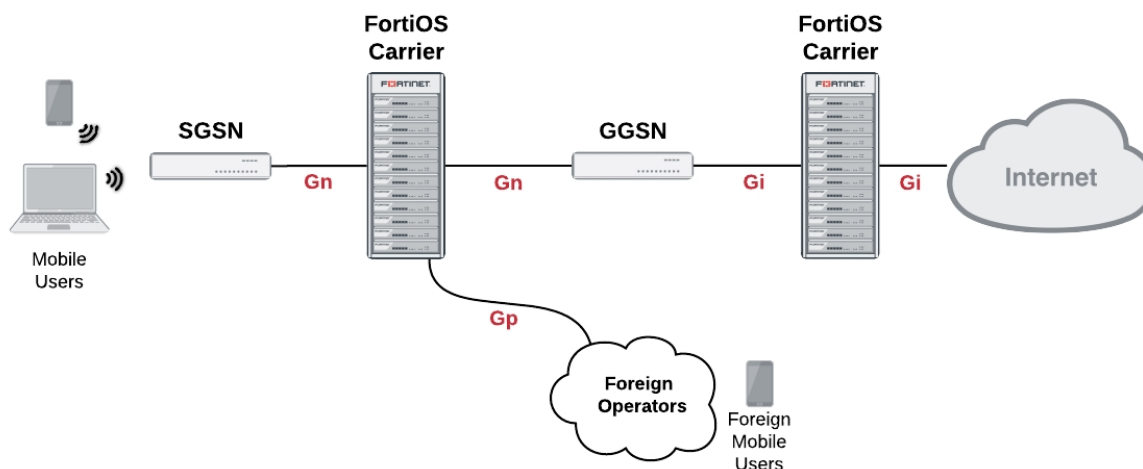
The parts of a GPRS network can be separated into the following groups according to the roles of the devices:

- Radio access to the GPRS network is accomplished by mobile phones and mobile stations (MS).
- Transport the GPRS packets across the GPRS network is accomplished by SGSNs and GGSNs, both local and remote, by delivering packets to the external services.
- Billing and records are handled by CDF, CFR, HLR, and VLR devices.

GPRS networks also rely on access points and PDP contexts as central parts of the communication structure. These are not actual devices, but they are still critical .

These devices, their roles, neighboring devices, the interfaces and protocols they use are outlined in the following table.

Carrier network showing the interfaces used (GTPv1)



Devices on a GTPv1 network

Device role	Neighboring Devices	Interfaces used	Protocols used
Mobile Users	Mobile Stations (MS)	Radio Access Technology (RAT)	

Device role	Neighboring Devices	Interfaces used	Protocols used
Mobile Stations (MS)	Mobile Users, SGSN	Gb	IP, Frame Relay
SGSN (local)	MS, SGSN (local or remote), GGSN (local and remote), CDR, CFR, HLR, VLR	Ga, Gb, Gn, Gp, Gz	IP, Frame Relay, GTP, GTP'
SGSN (remote)	SGSN (local)	Gn	GTP
GGSN (local)	SGSN (local or remote), GGSN (local and remote), CDR, CFR, HLR, VLR	Ga, Gi, Gn, Gp, Gz	IP, GTP, GTP'
GGSN (remote)	SGSN (local), WAP gateway, Internet, other external services	Gi, Gp	IP, GTPv1
CDR, CFR	SGSN (local), GGSN (local)	Ga, Gz	GTP'
HLR, VLR	SGSN (local), GGSN (local)	Ga, Gz	GTP'

Radio access

For a mobile phone to access the GPRS core network, it must first connect to a mobile station. This is a cellular tower that is connected to the carrier network.

How the mobile phone connects to the mobile station (MS) is determined by what Radio Access Technologies (RATs) are supported by the MS.

Transport

Transport protocols move data along the carrier network between radio access and the Internet or other carrier networks.

FortiOS Carrier should be present where information enters the Carrier network, to ensure the information entering is correct and not malicious. This means a Carrier-enabled FortiGate unit intercepts the data coming from the SGSN or foreign networks destined for the SSGN or GGSN onto the network, and after the GGSN as the data is leaving the network.

GTP

GPRS Tunnelling Protocol (GTP) is a group of IP-based communications protocols used to carry General Packet Radio Service (GPRS) within Global System for Mobile Communications (GSM) and Universal Mobile Telecommunications System (UMTS) networks. It allows carriers to transport actual cellular packets over a network via tunneling. This tunneling allows users to move between SGSNs and still maintain connection to the Internet through the GGSN.

GTP has three versions version 0, 1, and 2. GTP1 and GTP2 are supported by FortiOS Carrier. The only GTP commands that are common to all forms of GTP are the echo request/response commands that allow GSNs to verify up to once every 60 seconds that neighboring GSNs are alive.

GTPv0

There have been three versions of GTP to date. The original version of GTP (version 0) has the following differences from version GTPv1.

- the tunnel identification is not random
- there are options for transporting X.25
- the fixed port number 3386 is used for all functions, not just charging
- optionally TCP is allowed as a transport instead of UDP
- not all message types are supported in version 0

GTPv1

On a GPRS network, Packet Data Protocol (PDP) context is a data structure used by both the Serving GPRS Support Node (SGSN) and the Gateway GPRS Support Node (GGSN). The PDP context contains the subscribers information including their access point, IP address, IMSI number, and their tunnel endpoint ID for each of the SGSN and GGSN.

The Serving GPRS Support Node (SGSN) is responsible for the delivery of data packets from and to the mobile stations within its geographical service area. Its tasks include packet routing and transfer, mobility management (attach/detach and location management), logical link management, and authentication and charging functions. The location register of the SGSN stores location information (e.g., current cell, current VLR) and user profiles (e.g., IMSI, address(es) used in the packet data network) of all GPRS users registered with this SGSN.

GTPv1-C

GTPv1-C refers to the control layer of the GPRS Transmission network. This part of the protocol deals with network related traffic.

FortiOS Carrier handles GTPv1-C in GTPv1 by using the Tunnel Endpoint Identifier (TEID), IP address and a Network layer Service Access Point Identifier (NSAPI), sometimes called the application identifier, as an integer value that is part of the PDP context header information used to identify a unique PDP context in a mobile station, and SGSN.

For more information on GTPv1-C, see GTP-C messages.

GTPv1-U

GTPv1-U is defined in 3GPP TS 29.281 and refers to the user layer of the GPRS Tunneling network. This part of the protocol deals with user related traffic, user tunnels, and user administration issues.

A GTPv1-U tunnel is identified by a TEID, an IP address, and a UDP port number. This information uniquely identifies the limb of a GTPv1 PDP context. The IP address and the UDP port number define a UDP/IP path, a connectionless path between two endpoints (i.e. SGSN or GGSN). The TEID identifies the tunnel endpoint in the receiving GTPv1-U protocol entity; it allows for the multiplexing and demultiplexing of GTP tunnels on a UDP/IP path between a given GSN-GSN pair. For more information on GTPv1-U, see GTP-U messages.

The GTP core network consists of one or more SGSNs and GGSNs.

GGSN

The Gateway GPRS Support Node (GGSN) connects the GPRS network on one side via the SGSN to outside networks such as the Internet. These outside networks are called packet data networks (PDNs). The GGSN acts as an edge router between the two different networks — the GGSN forwards incoming packets from the external PDN to the addressed SGSN and the GGSN also forwards outgoing packets to the external PDN. The GGSN also converts the packets from the GPRS packets with SGSN to the external packets, such as IP or X.25.


SGSN

The Serving GPRS Support Node (SGSN) connects the GPRS network to GTPv1 compatible mobile stations, and mobile units (such as UTRAN and ETRAN) on one side and to the gateway node (GGSN), which leads to external networks, on the other side. Each SGSN has a geographical area, and mobile phones in that area connect to the GPRS network through this SGSN. The SGSN also maintains a location register that contains customer's location and user profiles until they connect through a different SGSN at which time the customer information is moved to the new SGSN. This information is used for packet routing and transfer, mobility management also known as location management, logical link management, and authentication and billing functions.

GTPv2

GTPv2, defined in 3GPP TS 29.274, is dramatically different from GTPv1, defined in 3GPP TS 29.060. Where in GTPv1 the tunnel is between the SGSN and the GGSN, in GTPv2 The SGSN is between the MME and the LTE Serving Gateway (S-GW), beyond which is the PDN gateway (P-GW). Even tunnel management messages have changed significantly.

Network diagram for GTPv2



Device roles on a GTPv2 network

Device role	Neighboring Devices	Interfaces used	Protocols used
Mobile Users	Mobile Stations (MS)	Radio Access Technology (RAT)	--
GTPv1 Mobile Stations (MS)	Mobile Users, SGSN	Gb	IP, Frame Relay

Device role	Neighboring Devices	Interfaces used	Protocols used
GTPv2 Mobile Stations (MS)	Mobile Users, MME	???	IP, Frame Relay
SGSN (local)	GTPv1 MS, SGSN, S-GW	???	IP, Frame Relay, GTPv1, GTP'
S-GW	SGSN, MME, P-GW	???	IP, GTPv2, GTP'
P-GW	S-GW, Internet, other external services	???	IP, GTPv2

GTPv2-C

GTPv2-C is the control layer messaging for GTPv2. It is used by LTE mobile stations, SGSN units for backwards compatibility, and SGWs that are the gateway to other networks. The messaging is very different from GTPv1. GTPv2-C is required to communicate with the Mobility Management Entity (MME) to create, change and delete EPS bearers when handover events happen, and to create Forwarding tunnels. The protocol is also used to communicate with the Serving Gateway (SGW) which has the S-GW and PDN-GW interfaces, and the Serving GPRS Support Node (SGSN).

MME

MME essentially fills the role of the SGSN in a GTPv1 network — it is how the mobile stations gain access to the Carrier network. GTPv2 supports different mobile stations than GTPv1, so MME handles the GTPv2 MSes and SGSN handles the GTPv1 MSes

Billing and records

A major part of the GPRS network is devoted to billing. Customer billing requires enough information to identify the customer, and then billing specific information such as connection locations and times, as well as amount of data transferred. A modified form of GTP called GTP' is used for billing. The home location records and visitor location records store information about customers that is critical to billing.

GTP' (GTP prime)

GTP is used to handle tunnels of user traffic between SGSNs and GGSNs. However for billing purposes, other devices that are not supported by GTP are required. GTP' (GTP prime) is a modified form of GTP and is used to communicate with these devices such as the Charging Data Function (CDF) that communicates billing information to the Charging Gateway Function (CGF). In most cases, GTP' transports user records from many individual network elements, such as the GGSNs, to a centralism computer which then delivers the charging data more conveniently to the network operator's billing center, often through the CGF. The core network sends charging information to the CGF, typically including PDP context activation times and the quantity of data which the end user has transferred.

GTP' is used by the Ga and Gz interfaces to transfer billing information. GTP' uses registered UDP/TCP port 3386. GTP' defines a different header, additional messages, field values, as well as a synchronization protocol to avoid losing or duplicating CDRs on CGF or SGSN/GGSN failure. Transferred CDRs are encoded in ASN.1.

HLR

The Home Location Register (HLR) is a central database that contains details of each mobile phone subscriber that is authorized to use the GSM core network. There can be several logical, and physical, HLRs per public land mobile network (PLMN), though one international mobile subscriber identity (IMSI)/MSISDN pair can be associated with only one logical HLR (which can span several physical nodes) at a time. The HLRs store details of every SIM card issued by the mobile phone operator. Each SIM has a unique identifier called an IMSI which is the primary key to each HLR record.

VLR

The Visitor Location Register (VLR) is a database which stores information about all the mobile devices that are currently under the jurisdiction of the Mobile Switching Center which it serves. Of all the information the VLR stores about each Mobile Station, the most important is the current Location Area Identity (LAI). This information is vital in the call setup process.

Whenever an MSC detects a new MS in its network, in addition to creating a new record in the VLR, it also updates the HLR of the mobile subscriber, informing it of the new location of that MS.

For more information on GTP', see GTP-U and Charging Management Messages.

GPRS network common interfaces

There are interfaces for each connection on the GPRS network. An interface is an established standard form of communication between two devices. Consider a TCP/IP network. In addition to the transport protocol (TCP) there are other protocols on that network that describe how devices can expect communications to be organized, just like GPRS interfaces.

Interfaces between devices on the network

There are a series of interfaces that define how different devices on the carrier network communicate with each other. These interfaces are called Ga to Gz, and each one defines how a specific pair of devices will communicate. For example Gb is the interface between the base station and the SGSN, and Gn is one possible interface between the SGSN and GGSN.

The SGSN and GGSN keep track of the CDR information and forward it to the Charging Data Function (CDF) using the Gr interface between the SGSN and home location register (HLR), Gs interface between the SGSN and MSC (VLR), Gx interface between the GGSN and the Charging Rules Function (CRF), Gy between the GGSN and online charging system (OCS), and finally Gz which is the off-line (CDR-based) charging interface between the GSN and the CG that uses GTP'.

Each of these interfaces on the GPRS network has a name in the format of G_x where x is a letter of the alphabet that determines what part of the network the interface is used in. It is common for network diagrams of GPRS networks to include the interface name on connections between devices.



The Carrier-enabled FortiGate unit only provides protection on the Gn, Gp, and Gi interfaces.

GPRS network interfaces, their roles, and billing

Name	Device connections that use this interface	Traffic Protocol used	Its role or how it affects billing
Ga	CDR and GSN (SGSNs and GGSNs)	GTP ^a - GTP modified to include CDR role	CDR have the accounting records, that are compiled in the GSN and then sent to the Charging Gateway (CG)
Gb	MS and SGSN	Frame Relay or IP	When an IP address moves to a new MS, the old MS may continue to use and bill that IP address.
Gi	GGSN and public data networks (PDNs)	IP based	This is the connection to the Internet. If the GTP tunnel is deleted without notifying the Gi interface, the connection may remain open incurring additional charges. FortiOS Carrier adds this interface to a firewall. See Anti-overbilling with FortiOS Carrier.
Gn	SGSN and external SGSNs and internal GGSNs	GTP	When the GTP tunnel is deleted, need to inform other interfaces immediately to prevent misuse of connections remaining open. FortiOS Carrier adds this interface to a firewall.
Gp	Internal SGSN and external GGSNs	GTP	
Gz	GSN (SGSN and GGSN) and the charging gateway (CG)	GTP ^a	Used for the offline charging interface. Ga is used for online charging.

Corporate customers may have a direct connection to the Gi interface for higher security. The Gi interface is normally an IP network, though a tunnelling protocol such as GRE or IPsec may be used instead.

GTP Configuration

The GTP (GPRS Tunneling Protocol) is one of the major mobile core protocols used since to transfer data in the core mobile network. Mobility and data are exploding and this trend will continue with VoLTE, 5G, and the Internet of Things (IoT). The role of GTP in mobile networks will continue to remain critical.

With the mobile network ever growing importance as the communication channel for data rich application on mobile devices, connected intelligent devices and the IoT, comes the growing potential for attacks on the mobile infrastructure.

Introduction to GTP

GTP as a Potential Attack Vector

GTP's role in transferring data in the core mobile infrastructure makes it a potential ideal attack vector. To understand the security features for GTP we need to understand the risks that might compromise this protocol.

The business impact might vary in-between the different attacks from Denial of Service (DoS) attacks that hinders the capability of performing a legitimate operation due to resource starvation (for example - not being able to charge the customer for GPRS traffic use due to denial of service attack on the Charging GW) to remote compromise attacks that allows the hacker to have remote control of a critical device (for example – take control over a GGSN).

GTP-based attacks may have a wide range of business impact, based on the attacked devices' vulnerability, ranging from service unavailability, compromise customer information, and gaining control over infrastructure elements, just to give a few examples.

Listed below are the main categories of GTP-based attacks:

- **Protocol anomaly attacks** are packets and packets formats that should not be expected on the GTP protocol. These can include malformed packets, reserved packets' fields and types, etc.
- **Infrastructure attacks** are attempts to connect to restricted core elements, such as the GGSN, SGSN, PGW, etc.
- **Overbilling attacks** results in customers charged for traffic they did not use or the opposite of not paying for the used traffic.

Protecting Against GTP-Based Attacks: The Carrier Grade GTP Firewall

With the evolution of the mobile network so has GTP evolved. The awareness to the potential of GTP-based attacks has led mobile core vendors to harden their software to better deal with a potential attack. Alongside this evolution, network security vendors, such as Fortinet, has led the way in providing specific GTP aware firewalls to secure and protect the different versions of the GTP protocol from potential attacks.

A GTP firewall should be placed where GTP traffic and session originate and terminate, as shown in the below diagram, and has to inspect both the GTP-C (Control Plane) and GTP-U (Data Plane) packets that, together, constitute the GPRS Tunneling Protocol.

The GTP firewall in both cases is placed in line between the SGSN / SGW and the GGSN / PGW which are the initiator and terminator of the GTP traffic. One of the main roles of GTP firewall is also to be able to support the roaming between different versions of GTP without interrupting the service.

The GTP firewall must be carrier grade in its ability to scale and provide high availability without impact its ability to provide effective protection.

FortiGate with FortiCarrier – The Leading GTP Firewall

FortiGate is Fortinet's physical security platform, built specifically for high performance and scalability with the utilization of specialized FortiASIC technology. Fortinet Content Processors (CP) and Network Processors (NP) enable, offloading CPU intensive tasks and allowing the FortiGate to provide carrier grade performance and scalability. Utilizing the power of the FortiGate platform, FortiOS, Fortinet's security Operating System, provides threat intelligence and advanced functionalities to provide effective security, ranging from Carrier Grade NAT (CGNAT), firewalling, IPSec, etc.

FortiCarrier is the part of FortiOS which was specifically designed to provide security for specific carriers and mobile operators' protocols and requirements, such as awareness and security for GTP. The wide range of FortiGate platforms with FortiOS and FortiCarrier enables mobile operators to cost effectively secure their mobile network against GTP-based attacks, while ensuring unparalleled performance, availability and security effectiveness.

GTP Profile

You can configure multiple GTP profiles within the GTP menu. GTP profiles concern GTP activity flowing through the unit. These GTP profiles are then applied to a security policy.

GTP profile configuration settings

The following are GTP profile configuration settings in **Security Profiles > GTP Profiles**.

GTP Profile	
Lists each GTP profile that you have created. On this page, you can edit, delete or create a new GTP profile.	
Create New	Creates a new GTP profile. When you select Create New , you are automatically redirected to the New page.
Edit	Modifies settings within a GTP profile in the list. When you select Edit , you are automatically redirected to Edit page.
Delete	Removes a GTP profile from the list.
	To remove multiple GTP profiles from within the list, on the GTP Profile page, in each of the rows of the profiles you want removed, select the check box and then select Delete .
	To remove all GTP profiles from within the list, on the GTP Profile page, select the check box in the check box column and then select Delete .
Name	The name of the GTP profile.

Displays the number of times the object is referenced to other objects. For example, av_1 profile is applied to a security policy; on the Profile page (**Security Profiles > Antivirus > Profiles**), 1 appears in **Ref.** .

To view the location of the referenced object, select the number in **Ref.**, and the Object Usage window appears displaying the various locations of the referenced object.

To view more information about how the object is being used, use one of the following icons that is available within the Object Usage window:

Ref.

- **View the list page for these objects** – automatically redirects you to the list page where the object is referenced at.
- **Edit this object** – modifies settings within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy and so, when this icon is selected, the user is redirected to the Edit Policy page.
- **View the details for this object** – table, similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with. For example, av_1 profile is referenced with a security policy, and that security policy's settings appear within the table.

New GTP Profile

Provides settings for configuring a GTP profile.

Name	Enter a name for the GTP profile.
General Settings	Configure general options for the GTP profile.
Message Type Filtering	Configure filtering for messages.
APN Filtering	Configure filtering options for APN.
Basic Filtering	Configure filtering options for IMSI.
Advanced Filtering	Configure advanced filtering options.
IE removal policy	Configure IE removal policy options.
Encapsulated IP Traffic Filtering	Configure filtering options for encapsulated IP traffic.
Encapsulated Non-IP End User Address Filtering	Configure filtering options for encapsulated non-IP end user addresses.
Protocol Anomaly	Configure protocol anomaly options.

Anti-Overbilling	Configure anti-overbilling options.
-------------------------	-------------------------------------

Log	Configure log options.
------------	------------------------

General settings options

The following are mostly house keeping options that appear in the General Settings area of the GTP configuration page.

General Settings section of the New GTP Profile

GTP-in-GTP	<p>Select Allow to enable GTP packets to be allowed to contain GTP packets, or a GTP tunnel inside another GTP tunnel.</p> <p>To block all GTP-in-GTP packets, select Deny.</p>
Minimum Message Length	<p>Enter the shortest possible message length in bytes. Normally this is controlled by the protocol, and will vary for different message types. If a packet is smaller than this limit, it is discarded as it is likely malformed and a potential security risk.</p> <p>The default minimum message length is 0 bytes.</p>
Maximum Message Length	<p>Enter the maximum allowed length of a GTP packet in bytes.</p> <p>A GTP packet contains three headers and corresponding parts GTP, UDP, and IP. If a packet is larger than the maximum transmission unit (MTU) size, it is fragmented to be delivered in multiple packets. This is inefficient, resource intensive, and may cause problems with some applications.</p> <p>By default the maximum message length is 1452 bytes.</p>
Tunnel Limit	<p>Enter the maximum number of tunnels allowed open at one time. For additional GTP tunnels to be opened, existing tunnels must first be closed.</p> <p>This feature can help prevent a form of denial of service attack on your network. This attack involves opening more tunnels than the network can handle and consuming all the network resources doing so. By limiting the number of tunnels at any one time, this form of attack will be avoided.</p> <p>The tunnel limiting applies to the Handover Group, and Authorized SGSNs and GGSNs.</p>

General Settings section of the New GTP Profile

Tunnel Timeout	<p>Enter the maximum number of seconds that a GTP tunnel is allowed to remain active. After the timeout the unit deletes GTP tunnels that have stopped processing data. A GTP tunnel may hang for various reasons. For example, during the GTP tunnel tear-down stage, the "delete pdap context response" message may get lost. By setting a timeout value, you can configure the FortiOS Carrier firewall to remove the hanging tunnels.</p> <p>The default is 86400 seconds, or 24 hours.</p>
Control plane message rate limit	<p>Enter the number of packets per second to limit the traffic rate to protect the GSNs from possible Denial of Service (DoS) attacks. The default limit of 0 does not limit the message rate.</p> <p>GTP DoS attacks can include:</p> <ul style="list-style-type: none"> • Border gateway bandwidth saturation: A malicious operator can connect to your GRX and generate high traffic towards your Border Gateway to consume all the bandwidth. • GTP flood: A GSN can be flooded by illegitimate traffic
Handover Group	<p>Select the allowed list of IP addresses allowed to take over a GTP session when the mobile device moves locations.</p> <p>Handover is a fundamental feature of GPRS/UMTS, which enables subscribers to seamlessly move from one area of coverage to another with no interruption of active sessions. Session hijacking can come from the SGSN or the GGSN, where a fraudulent GSN can intercept another GSN and redirect traffic to it. This can be exploited to hijack GTP tunnels or cause a denial of service.</p> <p>When the handover group is defined it acts like a white list with an implicit default deny at the end — the GTP address must be in the group or the GTP message will be blocked. This stops handover requests from untrusted GSNs.</p>
Authorized SGSNs	<p>Use Authorized SGSNs to only allow authorized SGSNs to send packets through the unit and to block unauthorized SGSNs. Go to Firewall Objects > Address > Addresses and add the IP addresses of the authorized SGSNs to a firewall address or address group. Then set Authorized SGSNs to this firewall address or address group.</p> <p>You can use Authorized SGSNs to allow packets from SGSNs that have a roaming agreement with your organization.</p>

General Settings section of the New GTP Profile

Authorized GGSNs

Use **Authorized GGSNs** to only allow authorized GGSNs to send packets through the unit and to block unauthorized GGSNs. Go to **Firewall Objects > Address > Addresses** and add the IP addresses of the authorized GGSNs to a firewall address or address group. Then set **Authorized GGSNs** to this firewall address or address group.

You can use Authorized GGSNs to allow packets from SGSNs that have a roaming agreement with your organization.

Message type filtering options

On the **New GTP Profile** page, you can select to allow or deny the different types of GTP messages, which is referred to as message type filtering. You must expand the Message Type Filtering section to access the settings.

The messages types include Path Management, Tunnel Management, Location Management, Mobility Management, MBMS, and GTP-U and Charging Management messages.



For enhanced security, Fortinet best practices dictate that you set Unknown Message Action to deny. This will block all unknown GTP message types, some of which may be malicious.

To configure message type filter options, expand **Message Type Filtering** in the GTP profile.

APN filtering options

An Access Point Name (APN) is an Information Element (IE) included in the header of a GTP packet. It provides information on how to reach a network.

An APN has the following format:

```
<network_id>[.mnc<mnc_int>.mcc<mcc_int>.gprs]
```

Where:

- **<network_id>** is a network identifier or name that identifies the name of a network, for example, `example.com` or `internet`.
- **[.mnc<mnc_int>.mcc<mcc_int>.gprs]** is the optional operator identifier that uniquely identifies the operator's PLMN, for example `mnc123.mcc456.gprs`.

Combining these two examples results in a complete APN of `internet.mnc123.mcc456.gprs`.

By default, the unit permits all APNs. However, you can configure APN filtering to restrict roaming subscribers' access to external networks.

APN filtering applies only to the GTP **create pdp request** messages. The unit inspects GTP packets for both APN and selected modes. If both parameters match and APN filter entry, the unit applies the filter to the traffic.

Additionally, the unit can filter GTP packets based on the combination of an IMSI prefix and an APN.



You cannot add an APN when creating a new profile.

APN Filtering	
Enable APN Filter	Select to enable APN filtering.
Default APN Action	Select the default action for APN filtering. If you select Allow , all sessions are allowed except those blocked by individual APN filters. If you select Deny , all sessions are blocked except those allowed by individual APN filters.
Value	The APN to be filtered.
Mode	The type of mode chosen that indicates where the APN originated and whether the Home Location Register (HLR) has verified the user subscription:
Action	The type of action that will be taken.
Edit	Modifies the settings within the filter. When you select Edit , the Edit window appears, which allows you to modify the settings of the APN.
Delete	Removes the APN from the list within the table, in the APN Filtering section.
Add APN	Adds a new APN filter to the list. When you select Add APN , the New window appears, which allows you to configure the APN settings.
New APN page	
Value	Enter an APN to be filtered. You can include wild cards to match multiple APNs. For example, the value internet* would match all APNs that begin with internet .
Mode	Select one or more of the available modes to indicate where the APN originated and whether the Home Location Register (HLR) has verified the user subscription.
Mobile Station provided	MS-provided APN, subscription not verified, indicates that the mobile station (MS) provided the APN and that the HLR did not verify the user's subscription to the network.
Network provided	Network-provided APN, subscription not verified, indicates that the network provided a default APN because the MS did not specify one, and that the HLR did not verify the user's subscription to the network.
Subscription Verified	MS or Network-provided APN, subscription verified, indicates that the MS or the network provided the APN and that the HLR verified the user's subscription to the network.
Action	Select Allow or Deny .

Basic filtering options

The International Mobile Station Identity (IMSI) is used by a GPRS Support Node (GSN) to identify a mobile station. Three elements make up every IMSI:

- the mobile country code (MCC)
- the mobile network code (MNC)
- the mobile subscriber identification number (MSIN).

The subscriber's home network—the public land mobile network (PLMN)—is identified by the IMSI prefix, formed by combining the MCC and MNC.

By default, the unit allows all IMSIs. You can add IMSI prefixes to deny GTP traffic coming from non-roaming partners. Any GTP packets with IMSI prefixes not matching the prefixes you set will be dropped. GTP **Create pdp** request messages are filtered and only IMSI prefixes matching the ones you set are permitted. Each GTP profile can have up to 1000 IMSI prefixes set.

An IMSI prefix and an APN can be used together to filter GTP packets if you set an IMSI filter entry with a non-empty APN.



You cannot add an IMSI when creating a new profile. You must add it after the profile has been created and you are editing the profile.

IMSI Filtering section of the New GTP Profile

Enable IMSI Filter	Select to enable IMSI filtering.
Default IMSI Action	Select the default action for IMSI filtering. If you select Allow , all sessions are allowed except those blocked by individual IMSI filters. If you select Deny , all sessions are blocked except those allowed by individual IMSI filters.
APN	The APN that is part of the IMSI that will be filtered.
MCC-MNC	The MCC-MNC part of the IMSI that will be filtered.
Mode	The type of mode that indicates where the APN originated and whether the Home Location Register (HLR) has verified the user subscription.
Action	The type of action that will be taken.
Edit	Modifies settings to an IMSI filter. When you select Edit , the Edit window appears, which allows you to modify the IMSI filter's settings.
Delete	Removes an IMSI filter from within the table, in the IMSI Filtering section.
Add IMSI	Adds a new IMSI filter to the list. When you select Add IMSI , the New window appears, which allows you to configure IMSI filter settings.
New IMSI page	

APN	Enter the APN part of the IMSI to be filtered.
MCC-MNC	Enter the MCC-MCC part of the IMSI to be filtered.
Mode	Select one or more of the available modes to indicate where the APN originated and whether the Home Location Register (HLR) has verified the user subscription.
Mobile Station provided	MS-provided APN, subscription not verified, indicates that the mobile station (MS) provided the APN and that the HLR did not verify the user's subscription to the network.
Network provided	Network-provided APN, subscription not verified, indicates that the network provided a default APN because the MS did not specify one, and that the HLR did not verify the user's subscription to the network.
Subscription Verified	MS or Network-provided APN, subscription verified, indicates that the MS or the network provided the APN and that the HLR verified the user's subscription to the network.
Action	Select Allow or Deny .

Advanced filtering options

The FortiOS Carrier firewall supports advanced filtering against the attributes RAT, RAI, ULI, APN restriction, and IMEI-SV in GTP to block specific harmful GPRS traffic and GPRS roaming traffic. The following table shows some of the GTP context requests and responses that the firewall supports.

Attributes supported by FortiCarrier firewalls

	GTP Create PDP Context Request	GTP Create PDP Context Response	GTP Update PDP Context Request	GTP Update PDP Context Response
APN	yes	yes	-	
APN Restriction	yes	-	-	yes
IMEI-SV	yes	-	-	-
IMSI	yes	-	yes	-
RAI	yes	-	yes	-
RAT	yes	-	yes	-
ULI	yes	-	yes	-

When editing a GTP profile, select **Advanced Filtering > Create New** to create and add a rule. When the rule matches traffic it will either allow or deny that traffic as selected in the rule.

Advanced Filtering	
Enable	Select to enable advanced filtering.
Default Action	Select the default action for advanced filtering. If you select Allow , all sessions are allowed except those blocked by individual advanced filters. If you select Deny , all sessions are blocked except those allowed by individual advanced filters.
Messages	The messages, for example, Create PDP Context Request.
APN Restriction	The APN restriction.
RAT Type	The RAT types associated with that filter.
ULI	The ULI pattern.
RAI	The RAI pattern.
IMEI	The IMEI pattern.
Action	The action that will be taken.
Edit	Modifies the filter's settings. When you select Edit , the Edit window appears, which allows you to modify the filter's settings.
Delete	Removes a filter from the list.
Add	Adds a filter to the list. When you select Add , the New window appears, which allows you to configure settings for messages, APN, IMSI, MSISDN, RAT type, ULI, RAI, IMEI patterns as well as the type of action.
New Filtering page	
Messages	The PDP content messages this profile will match.
Create PDP Context Request	Select to allow create PDP context requests.
Create PDP Context Response	Select to allow create PDP context responses.
Update PDP Context Request	Select to allow update PDP context requests.

Update PDP Context Response	Select to allow update PDP context responses.
APN	Enter the APN.
APN Mode	<p>Select an APN mode as one or more of</p> <ul style="list-style-type: none"> • Mobile Station provided • Network provided • Subscription provided <p>This field is only available when an APN has been entered.</p>
Mobile Station provided	MS-provided PAN, subscription not verified, indicates that the mobile station (MS) provided the APN and that the HLR did not verify the user's subscription to the network.
Network provided	Network-provided APN, subscription not verified, indicates that the network provided a default APN because the MS did not specify one, and that the HLR did not verify the user's subscription to the network.
Subscription verified	MS or Network-provided APN, subscription verified, indicates that the MS or the network provided the APN and that the HLR verified the user's subscription to the network.
APN Restriction	<p>Select the type of restriction that you want. You can choose all of the types, or one of the types. You cannot choose multiple types. Types include:</p> <ul style="list-style-type: none"> • all • Public-1 • Public-2 • Private-1 • Private-2
IMSI	Enter the IMSI.
MSISDN	Enter the MSISDN.
RAT Type	<p>Optionally select the RAT type as any combination of the following:</p> <ul style="list-style-type: none"> • Any • UTRAN • GERAN • Wifi • GAN • HSPA <p>Some RAT types are GTPv1 specific.</p>

ULI pattern	Enter the ULI pattern.
RAI pattern	Enter the RAI pattern.
IMEI pattern	Enter the IMEI pattern.
Action	Select either Allow or Deny .

Adding an advanced filtering rule

When adding a rule, use the following formats:

- Prefix, for example, range 31* for MCC matches MCC from 310 to 319.
- Range, for example, range 310-319 for MCC matches MCC from 310 to 319.
- Mobile Country Code (MCC) consists of three digits. The MCC identifies the country of domicile of the mobile subscriber.
- Mobile Network Code (MNC) consists of two or three digits for GSM/UMTS applications. The MNC identifies the home PLMN of the mobile subscriber. The length of the MNC (two or three digits) depends on the value of the MCC. Best practices dictate not to mix two and three digit MNC codes within a single MCC area.
- Location Area Code (LAC) is a fixed length code (of 2 octets) identifying a location area within a PLMN. This part of the location area identification can be coded using a full hexadecimal representation except for the following reserved hexadecimal values: 0000 and FFFE. These reserved values are used in some special cases when no valid LAI exists in the MS (see 3GPP TS 24.008, 3GPP TS 31.102 and 3GPP TS 51.011).
- Routing Area Code (RAC) of a fixed length code (of 1 octet) identifies a routing area within a location.
- CI or SAC of a fixed length of 2 octets can be coded using a full hexadecimal expression.
- Type Allocation Code (TAC) has a length of 8 digits.
- Serial Number (SNR) is an individual serial number identifying each equipment within each TAC. SNR has a length of 6 digits.
- Software Version Number (SVN) identifies the software version number of the mobile equipment. SVN has a length of 2 digits.



You cannot add an advanced filtering rule when creating a new profile. You must add it after the profile has been created and you are editing the profile.

Information Element (IE) removal policy options

In some roaming scenarios, the unit is installed on the border of the PLMN and the GRX. In this configuration, the unit supports information element (IE) removal policies to remove any combination of R6 IEs (RAT, RAI, ULI, IMEI-SV and APN restrictions) from the types of messages described in “[Advanced filtering options](#)”, prior to forwarding the messages to the HGGSN (proxy mode).

IE removal policy	
Enable	Select to enable this option.

SGSN address of message IE	The firewall address or address group that contains the SGSN addresses.
IEs to be removed	The IE types that will be removed. These include APN Restriction, RAT, RAI, ULI, and IMEI.
Add	Adds an IE removal policy. When you select Add , the New window appears, which allows you to configure the IE policy.
Edit	Modifies settings from within the IE removal policy. When you select Edit , the Edit window appears, which allows you to modify the settings within the policy.
Delete	Removes the IE removal policy from the list.
New IE policy page	
SGSN address	Select a firewall address or address group that contains SGSN addresses.
IEs to be removed	Select one or more IE types to be removed. These include APN Restriction, RAT, RAI, ULI, and IMEI.

Encapsulated IP traffic filtering options

You can use encapsulated IP traffic filtering to filter GTP sessions based on information contained in the data stream. to control data flows within your infrastructure. You can configure IP filtering rules to filter encapsulated IP traffic from mobile stations by identifying the source and destination policies. For more information, see When to use encapsulated IP traffic filtering.

Expand **Encapsulated IP Traffic Filtering** in the GTP profile to reveal the options.

Encapsulated IP Traffic Filtering	
Enable IP Filter	Select to enable encapsulated IP traffic filtering options.
Default IP Action	Select the default action for encapsulated IP traffic filtering. If you select Allow , all sessions are allowed except those blocked by individual encapsulated IP traffic filters. If you select Deny , all sessions are blocked except those allowed by individual encapsulated IP traffic filters.
Source	Select a source IP address from the configured firewall IP address or address group lists. Any encapsulated traffic originating from this IP address will be a match if the destination also matches.
Destination	Select a destination IP address from the configured firewall IP address or address group lists. Any encapsulated traffic being sent to this IP address will be a match if the destination also matches.

Action	The type of action that will be taken. Select to Allow or Deny encapsulated traffic between this source and Destination.
Edit	Modifies the source, destination or action settings.
Add IP Policy	Adds a new encapsulated IP traffic filter. When you select Add IP Policy , the New window appears which allows you to configure IP policy settings.
New (window)	
Source	Select the source firewall address or address group.
Destination	Select the destination firewall address or address group.
Action	Select Allow or Deny .

Encapsulated non-IP end user traffic filtering options

Depending on the installed environment, it may be beneficial to detect GTP packets that encapsulate non-IP based protocols. You can configure the FortiOS Carrier firewall to permit a list of acceptable protocols, with all other protocols denied.

The encoded protocol is determined in the PDP Type Organization and PDP Type Number fields within the End User Address Information Element. The PDP Type Organization is a 4-bit field that determines if the protocol is part of the ETSI or IETF organizations. Values are zero and one, respectively. The PDP Type field is one byte long. Both GTP specifications list only PPP, with a PDP Type value of one, as a valid ETSI protocol. PDP Types for the IETF values are determined in the "Assigned PPP DLL Protocol Numbers" sections of RFC1700. The PDP types are compressed, meaning that the most significant byte is skipped, limiting the protocols listed from 0x00 to 0xFF.

Encapsulated Non-IP End User Address Filtering	
Enable Non-IP Filter	Select to enable encapsulated non-IP traffic filtering.
Default Non-IP Action	Select the default action for encapsulated non-IP traffic filtering. If you select Allow , all sessions are allowed except those blocked by individual encapsulated non-IP traffic filters. If you select Deny , all sessions are blocked except those allowed by individual encapsulated non-IP traffic filters.
Type	The type chosen, AESTI or IETF .
Start Protocol	The beginning protocol port number range.
End Protocol	The end of the protocol port number range.
Action	The type of action that will be taken.

Edit	Modify a non-IP filter's settings in the list. When you select Edit , the Edit window appears, which allows you to modify the Non-IP policy settings.
Delete	Remove a non-IP policy from the list.
Add Non-IP Policy	Add a new encapsulated non-IP traffic filter. When you select Add Non-IP Policy , you are automatically redirected to the New page.
New (window)	
Type	Select AESTI or IETF .
Start Protocol	Select a start and end protocol from the list of protocols in RFC 1700. Allowed range includes 0 to 255 (0x00 to 0xff). Some common protocols include: <ul style="list-style-type: none"> • 33 (0x0021) Internet Protocol • 35 (0x0023) OSI Network Layer • 63 (0x003f) NETBIOS Framing • 65 (0x0041) Cisco Systems • 79 (0x004f) IP6 Header Compression • 83 (0x0053) Encryption
End Protocol	
Action	Select Allow or Deny .

Protocol Anomaly prevention options

Use protocol anomaly detection options to detect or deny protocol anomalies according to GTP standards and tunnel state. Protocol anomaly attacks involve malformed or corrupt packets that typically fall outside of the protocol specifications. Packets cannot pass through if they fail the sanity check.

Protocol Anomaly	
Invalid Reserved Field	GTP version 0 (GSM 09.60) headers specify a number of fields that are marked as "Spare" and contain all ones (1). GTP packets that have different values in these fields are flagged as anomalies. GTP version 1 (GSM 29.060) makes better use of the header space and only has one, 1-bit, reserved field. In the first octet of the GTP version1 header, bit 4 is set to zero.
Reserved IE	Both versions of GTP allow up to 255 different Information Elements (IE). However, a number of Information Elements values are undefined or reserved. Packets with reserved or undefined values will be filtered.
Miss Mandatory IE	GTP packets with missing mandatory Information Elements (IE) will not be passed to the GGSN.

Out of State Message	<p>The GTP protocol requires a certain level of state to be kept by both the GGSN and SGSN. Some message types can only be sent when in a specific GTP state. Packets that do not make sense in the current state are filtered or rejected.</p> <p>Both versions of GTP allow up to 255 different message types. However, a number of message type values are undefined or reserved.</p> <p>Best practices dictate that packets with reserved or undefined values will be filtered.</p>
Out of State IE	GTP Packets with out of order Information Elements are discarded.
Spoofed Source Address	<p>The End User Address Information Element in the PDP Context Create & Response messages contain the address that the mobile station (MS) will use on the remote network. If the MS does not have an address, the SGSN will set the End User Address field to zero when sending the initial PDP Context Create message. The PDP Context Response packet from the GGSN will then contain an address to be assigned to the MS. In environments where static addresses are allowed, the MS will relay its address to the SGSN, which will include the address in the PDP Context Create Message. As the MS address is negotiated within the PDP Context creation handshake, any packets originating from the MS that contain a different source address are detected and dropped.</p>

Anti-Overbilling options

You can configure the FortiOS Carrier firewall to prevent over billing subscribers for traffic over the. To enable anti-overbilling, you must configure both the Gn/Gp firewall and the Gi firewall.

Expand **Anti-Overbilling** in the GTP profile to reveal these settings.

Anti-Overbilling	
Gi Firewall IP Address	The IP address of the unit's interface configured as a Gi gateway.
Port	The SG security port number. The default port number is port 21123. Change this number if your system uses a different SG port.
Interface	Select the unit interface configured as a Gi gateway.
Security Context ID	Enter the security context ID. This ID must match the ID entered on the server Gi firewall. The default security context ID is 696.

Log options

All the GTP logs are treated as a subtype of the event logs. To enable GTP logging, you must:

- configure the GTP log settings in a GTP profile

Log	
Log Frequency	<p>Enter the number of messages to drop between logged messages.</p> <p>An overflow of log messages can sometimes occur when logging rate-limited GTP packets exceed their defined threshold. To conserve resources on the syslog server and the Carrier-enabled FortiGate unit, you can specify that some log messages are dropped. For example, if you want only every twentieth message to be logged, set a logging frequency of 20. This way, 20 messages are skipped and the next logged.</p> <p>Acceptable frequency values range from 0 to 2147483674. When set to '0', no messages are skipped.</p>
Forwarded Log	Select to log forwarded GTP packets.
Denied Log	Select to log GTP packets denied or blocked by this GTP profile.
Rate Limited Log	Select to log rate-limited GTP packets.
State Invalid Log	Select to log GTP packets that have failed stateful inspection.
Tunnel Limit Log	Select to log packets dropped because the maximum limit of GTP tunnels for the destination GSN is reached.
Extension Log	<p>Select to log extended information about GTP packets. When enabled, this additional information will be included in log entries:</p> <ul style="list-style-type: none"> • IMSI • MSISDN • APN • Selection Mode • SGSN address for signaling • SGSN address for user data • GGSN address for signaling • GGSN address for user data

Traffic count Log

Select to log the total number of control and user data messages received from and forwarded to the GGSNs and SGSNs that the unit protects.

The unit can report the total number of user data and control messages received from and forwarded to the GGSNs and SGSNs it protects. Alternately, the total size of the user data and control messages can be reported in bytes. The unit differentiates between traffic carried by each GTP tunnel, and also between GTP-User and GTP-Control messages.

The number of messages or the number of bytes of data received from and forwarded to the SGSN or GGSN are totaled and logged if a tunnel is deleted.

When a tunnel is deleted, the log entry contains:

- Timestamp
- Interface name (if applicable)
- SGSN IP address
- GGSN IP address
- TID
- Tunnel duration time in seconds
- Number of messages sent to the SGSN
- Number of messages sent to the GGSN

Specifying logging types

You can configure the unit to log GTP packets based on their status with GTP traffic logging.

The status of a GTP packet can be any of the following 5 states:

- **Forwarded** - a packet that the unit transmits because the GTP policy allows it
- **Prohibited** - a packet that the unit drops because the GTP policy denies it
- **Rate-limited** - a packet that the unit drops because it exceeds the maximum rate limit of the destination GSN
- **State-invalid** - a packet that the unit drops because it failed stateful inspection
- **Tunnel-limited** - a packet that the unit drops because the maximum limit of GTP tunnels for the destination GSN is reached.

The following information is contained in each log entry:

- Timestamp
- Source IP address
- Destination IP address
- Tunnel Identifier (TID) or Tunnel Endpoint Identifier (TEID)
- Message type
- Packet status: forwarded, prohibited, state-invalid, rate-limited, or tunnel-limited
- Virtual domain ID or name
- Reason to be denied if applicable.

GTP performance

There are independent Receive and Transmit queues for `gtp-u` process. These queues are and their associated resources are initialized when the `ftp-enhance-mode` is enabled.

CLI changes under system npu

gtp-enhance-mode

```
config system npu
  set gtp-enhance-mode {enable|disable}
end
```



This configuration requires a reboot of the device to initialize the changes.

gtp-enhance-cpu-range

This is used to set the CPUs which can process the GTP-U packet inspection.

```
config system npu
  set gtp-enhance-cpu-range {0|1|2}
end
```

Option	Description
0	Inspect GTPU packets by all CPUs
1	Inspect GTPU packets by Master CPUs
2	Inspect GTPU packets by Slave CPUs

Diagnose commands

```
diagnose npu np6 hbq-stats {all|np xx}
```

Used to see the GTP-U packet counter by all NP or the corresponding np.

```
diagnose npu np6 hbq-stats-clear all /np xx
```

Used to clear the GTP-U packet counter by all NP or the corresponding np.

Verifying the enhance-mode is disabled

Before execute the test or enable/disable the gtp enhance, first check the `gtp-enhance-mode` status as in the example below:

```
config system npu
  get
  gtp-enhance-mode: disable
  gtp-enhance-cpu-range: 0
end
```

If the `gtp-enhance-mode` is `disable`, use the command `diagnose npu np6 hbq-stats all`.

The output will be similar to below:

```
# diagnose npu np6 hbq-stats all
Total :0
```

If the `gtp-enhance-mode` is enable, use the command `diagnose npu np6 hbq-stats all`

The output will be similar to below:

```
# diagnose npu np6 hbq-stats all
cpu_ 0:0
cpu_ 1:0
cpu_ 2:0
cpu_ 3:0
cpu_ 4:0
cpu_ 5:0
cpu_ 6:0
cpu_ 7:0
cpu_ 8:0
cpu_ 9:0
cpu_10:0
cpu_11:0
cpu_12:0
cpu_13:0
cpu_14:0
cpu_15:0
cpu_16:0
cpu_17:0
cpu_18:0
cpu_19:0
cpu_20:0
cpu_21:0
cpu_22:0
cpu_23:0
cpu_24:0
cpu_25:0
cpu_26:0
cpu_27:0
cpu_28:0
cpu_29:0
cpu_30:0
cpu_31:0
cpu_32:0
cpu_33:0
cpu_34:0
cpu_35:0
cpu_36:0
cpu_37:0
cpu_38:0
cpu_39:0
Total :0
```

Sometimes, when loading the new configure file, and the new configure file does not match the old configure file, the `gtp-enhance-mode` status will be confused.

You can see :

```
#config system npu
#get
gtp-enhance-mode: enable
```

but you can also see that

```
diagnose npu np6 hbq-stats all
Total :0
```

This means the `gtp-enhance-mode` is actually set to `disable`.

The inverse is also possible, when you see

```
#config system npu
#get
gtp-enhance-mode: disable
```

but you also see that

```
# diagnose npu np6 hbq-stats all
cpu_ 0:0
...
cpu_39:0
Total :0
```

This means the `gtp-enhance-mode` is actually set to `enable`.

If these combinations occur, just run the command below:

```
config system npu
  set gtp-enhance-mode enable
end
```

or

```
config system npu
  set gtp-enhance-mode disable
end
```

Once this is done, reboot the device to let the 2 statuses match.

Configuring GTP on FortiOS Carrier

Configuring GTP support on FortiOS Carrier involves configuring a number of areas of features.

GTP support on the Carrier-enabled FortiGate unit

The FortiCarrier unit needs to have access to all traffic entering and exiting the carrier network for scanning, filtering, and logging purposes. This promotes one of two configurations — hub and spoke, or bookend.

A hub and spoke configuration with the Carrier-enabled FortiGate unit at the hub and the other GPRS devices on the spokes is possible for smaller networks where a lower bandwidth allows you to divide one unit into multiple virtual domains to fill multiple roles on the carrier network. It can be difficult with a single FortiOS Carrier as the hub to ensure all possible entry points to the carrier network are properly protected from potential attacks such as relayed network attacks.

A bookend configuration uses two Carrier-enabled FortiGate units to protect the carrier network between them with high bandwidth traffic. One unit handles traffic from mobile stations, SGSNs, and foreign carriers. The other handles GGSN and data network traffic. Together they ensure the network is secure.

The Carrier-enabled FortiGate unit can access all traffic on the network. It can also verify traffic between devices, and verify that the proper GPRS interface is being used. For example there is no reason for a Gn interface to be used to communicate with a mobile station — the mobile station will not know what to do with the data — so that traffic is blocked.



When you are configuring your Carrier-enabled FortiGate unit's GTP profile, you must first configure the APN. It is critical to GTP communications — no traffic will flow without the APN.

The Carrier-enabled FortiGate unit does more than just forward and route GTP packets over the network. It also performs:

- [GTP support on the Carrier-enabled FortiGate unit](#)
- [GTP support on the Carrier-enabled FortiGate unit](#)
- [GTP support on the Carrier-enabled FortiGate unit](#)
- [GTP support on the Carrier-enabled FortiGate unit](#)
- [GTP support on the Carrier-enabled FortiGate unit](#)

Packet sanity checking

The FortiOS Carrier firewall checks the following items to determine if a packet confirms to the UDP and GTP standards:

- GTP release version number — must be 0, 1, or 2
- Settings of predefined bits
- Protocol type
- UDP packet length

If the packet in question does not confirm to the standards, the FortiOS Carrier firewall drops the packet, so that the malformed or forged traffic will not be processed.

GTP stateful inspection

Apart from the static inspection (checking the packet header), the FortiOS Carrier firewall performs stateful inspection.

Stateful inspection provides enhanced security by keeping track of communications sessions and packets over a period of time. Both incoming and outgoing packets are examined. Outgoing packets that request specific types of incoming packets are tracked; only those incoming packets constituting a proper response are allowed through the firewall.

The FortiOS Carrier firewall can also index the GTP tunnels to keep track of them.

Using the enhanced Carrier traffic policy, the FortiOS Carrier firewall can block unwanted encapsulated traffic in GTP tunnels, such as infrastructure attacks. Infrastructure attacks involve attempts by an attacker to connect to restricted machines, such as GSN devices, network management systems, or mobile stations. If these attempts to connect are detected, they are to be flagged immediately by the firewall .

Protocol anomaly detection and prevention

The FortiOS Carrier firewall detects and optionally drops protocol anomalies according to GTP standards and specific tunnel states. Protocol anomaly attacks involve malformed or corrupt packets that typically fall outside of protocol specifications. These packets are not seen on a production network. Protocol anomaly attacks exploit poor programming practices when decoding packets, and are typically used to maliciously impair system performance or elevate privileges.

FortiOS Carrier also detects IP address spoofing inside GTP data channel.

See [Protocol anomaly detection and prevention](#).

GTP Shared Tunnel Limit

The GTP Shared Tunnel limit is the total number of GTP tunnels created by multiple GTP profiles. A Global shared tunnel limit gives the flexibility of limiting the number of GTP tunnels flowing through different profiles. A shared global limit is defined and then referenced in the profiles. Before FortiOS 6.0, the GTP tunnel limit could be set on a per-profile basis. The GTP tunnel limits can now be set per VDOM.

Per profile tunnel limiting is still possible but restrictive limits between the global limit and the profile limit will be enforced.

Example:

- Global shared tunnel limit defined as 12
- GTP Profile A - per profile tunnel limit defined as 8
- GTP Profile B - per profile tunnel limit defined as 14

You can have eight tunnels active in profile A, but the ninth will be dropped due to the profile limit of 8.

If Profile A still has eight active tunnels, you can have four tunnels active in profile B and the fifth will be dropped even though the profile allows fourteen, because the global share limit is twelve.

Configuring the GTP tunnel limit

The tunnel limit shaper is created/edited by configuring the object `config gtp tunnel-limit`.

Each GTP profile selects a shaper to use from the object configured in the CLI.

These values are shared by all of the profiles associated with the indicated shaper:

- Each GTP profile selects the shared tunnel limiter from the CLI.
- Each shared tunnel limiter counts the total number of alive GTP tunnels created by GTP profiles that select the limiter.
- Once the total number exceeds the shared limits, new tunnel requests of those GTP profiles will be rejected.

CLI syntax:

```
config gtp tunnel-limit
  edit "gtp-tl-1"
    set tunnel-limit <integer value from 1 to 16000000>
  end
```

Use the syntax below to assign the tunnel limit to the GTP profile.

CLI syntax:

```
config firewall gtp
  edit "gtp1"
    set global-tunnel-limit "gtp-tl-1"
  next
  edit "gtp2"
    set global-tunnel-limit "gtp-tl-1"
  end
```

Diagnose command

A diagnose command is available to show the shared tunnel limiters.

```
diagnose firewall gtp tunnel-limit list
```

Example output:

```
name=gtp-tl-1 tunnel_limit=50 tunnel_count=0
```

HA

FortiOS Carrier active-passive HA provides failover protection for the GTP tunnels. This means that an active-passive cluster can provide FortiOS Carrier firewall services even when one of the cluster units encounters a problem that would result in complete loss of connectivity for a stand-alone FortiOS Carrier firewall. This failover protection provides a backup mechanism that can be used to reduce the risk of unexpected downtime, especially for mission-critical environments.

FortiOS HA synchs TCP sessions by default, but UDP sessions are not synchronized by default. However synchronizing a session is only part of the solution if the goal is to continue GTP processing on a synchronized session after a HA switch. For that to be successful we also need to synch the GTP tunnel state. So, once the master completes tunnel setup then the GTP tunnel is synchronized to the slave.

GTP traffic will only flow without interruption on a HA switch if bidirectional GTP policies have been configured: an internal (GTP server) to external (all) UDP port GTP policy, and an external (all) to internal (GTP server) UDP port GTP policy. If either policy is missing then traffic may be interrupted until traffic flows in the opposite direction.

For more information on HA in FortiOS, see the High Availability (HA) Guide or the FortiOS Administration Guide.

Virtual domain support

FortiOS Carrier is suited to both large and smaller carriers. A single Carrier-enabled FortiGate unit can serve either one large carrier, or several smaller ones through virtual domains. As with any FortiGate unit, Carrier-enabled units have the ability to split their resources into multiple virtual units. This allows smaller carriers to use just the resources that they need without wasting the extra. For more information on HA in FortiOS, see the Virtual Domains (VDOMs) Guide.

Configuring General Settings on the Carrier-enabled FortiGate unit

To configure the GTP General Settings, go to **Security Profiles > GTP Profile**, and edit a GTP profile. Expand **General Settings** to configure settings. See General settings options.

GTP Monitor Mode

The `monitor-mode` setting is part of the GTP profile. The setting shows on all GTP profiles and works for all GTP versions.

When this setting is enabled, if a GTP packet is to be dropped due to a GTP deny case such as:

- GTP_DENY
- GTP_RATE_LIMIT
- GTP_STATE_INVALID
- GTP_TUNNEL_LIMIT

instead of being dropped, it will be forwarded and logged with the original deny log message and a "-monitor" suffix (e.g., state-invalid-monitor).

This setting is found in the CLI.

```
config firewall gtp
edit profile_name
...
set monitor-mode {disable*|enable}
...
end
end
```

GTP Stats via SNMP

All parameters/values that can be checked with the following debug commands are available via SNMP:

- `diagnose firewall gtp stat`
- `diagnose firewall gtp runtime-stat`

The following OID was added to make GTP stats available via SNMP:

- 1.3.6.1.4.1.12356.101.5.3

The sequence of printed items of the debug commands was also adjusted to make it consistent with the data type and SNMP.

Configuring Encapsulated Filtering in FortiOS Carrier

Encapsulated traffic on the GPRS network can come in a number of forms as it includes traffic that is “wrapped up” in another protocol. This detail is important for firewalls because it requires “unwrapping” to properly scan the data inside. If encapsulated packets are treated as regular packets, that inside layer will never be scanned and may allow malicious data into your network.

On Carrier-enabled FortiGate units, GTP related encapsulated filtering falls under encapsulated IP traffic filtering, and encapsulated non-IP end user address filtering.

Configuring Encapsulated IP Traffic Filtering

Generally there are a very limited number of IP addresses that are allowed to encapsulate GPRS traffic. For example GTP tunnels are a valid type of encapsulation when used properly. This is the GTP tunnel which uses the Gp or Gn interfaces between SGSNs and GGSNs. However, a GTP tunnel within a GTP tunnel is not accessible — FortiOS Carrier will either block or forward the traffic, but is not able to open it for inspection.

The ability to filter GTP sessions is based on information contained in the data stream and provides operators with a powerful mechanism to control data flows within their infrastructure. You can also configure IP filtering rules to filter encapsulated IP traffic from Mobile Stations.

To configure the Encapsulated IP Traffic Filtering, go to **Security Profiles > GTP Profile**, and edit a GTP profile. Expand **Encapsulated IP Traffic Filtering** to configure settings. See Encapsulated IP traffic filtering options.

When to use encapsulated IP traffic filtering

The following are the typical cases that need encapsulated IP traffic filtering:

Mobile station IP pools

In a well-designed network, best practices dictate that the mobile station address pool is to be completely separate from the GPRS network infrastructure range of addresses. Encapsulated IP packets originating from a mobile station will not contain source or destination addresses that fall within the address range of GPRS infrastructures. In addition, traffic originating from the users handset will not have destination/source IP addresses that fall within any Network Management System (NMS) or Charging Gateway (CG) networks.

Communication between mobile stations

Mobile stations on the same GPRS network are not able to communicate with other mobile stations. Best practices dictate that packets containing both source and destination addresses within the mobile station's range of addresses are to be dropped.

Direct mobile device or internet attacks

It may be possible for attackers to wrap attack traffic in GTP protocols and submit the resulting GTP traffic directly to a GPRS network element from their mobile stations or a node on the Internet. It is possible that the receiving SGSN or GGSN would then strip off the GTP header and attempt to route the underlying attack. This underlying attack could have any destination address and would probably have a source address spoofed as if it were valid from that PLMN.



You cannot add an IE removal policy when you are creating a new profile.

Relayed network attacks

Depending on the destination the attack could be directly routed, such as to another node of the PLMN, or re wrapped in GTP for transmission to any destination on the Internet outside the PLMN depending on the routing table of the GSN enlisted as the unwitting relay.

The relayed attack could have any source or destination addresses and could be any of numerous IP network attacks, such as an attack to hijack a PDP context, or a direct attack against a management interface of a GSN or other device within the PLMN. Best practices dictate that any IP traffic originating on the Internet or from an MS with a destination address within the PLMN is to be filtered.

Configuring Encapsulated Non-IP End User Address Filtering

Much of the traffic on the GPRS network is in the form of IP traffic. However some parts of the network do not use IP based addressing, so the Carrier-enabled FortiGate unit is unable to perform Encapsulated IP Traffic Filtering.

Depending on the installed environment, it may be beneficial to detect GTP packets that encapsulate non-IP based protocols. You can configure the FortiOS Carrier firewall to permit a list of acceptable protocols, with all other protocols denied.

The encoded protocol is determined in the PDP Type Organization and PDP Type Number fields within the End User Address Information Element. The PDP Type Organization is a 4-bit field that determines if the protocol is part of the ETSI or IETF organizations. Values are zero and one, respectively. The PDP Type field is one byte long. Both GTP specifications only list PPP, with a PDP Type value of one, as a valid ETSI protocol. PDP Types for the IETF values are determined in the "Assigned PPP DLL Protocol Numbers" sections of RFC 1700. The PDP

types are compressed, meaning that the most significant byte is skipped, limiting the protocols listed from 0x00 to 0xFF.

To configure the Encapsulated Non-IP End User Address Filtering, go to **Security Profiles > GTP Profile**, and edit a GTP profile. Expand **Encapsulated Non-IP End User Address Filtering** to configure settings. See Encapsulated non-IP end user traffic filtering options.

Configuring the Protocol Anomaly feature in FortiOS Carrier

When anomalies do happen, it is possible for the anomaly to interrupt network traffic or consume network resources — if precautions are not taken. Anomalies can be generated by accident or maliciously, but both methods can have the same results — degrading the performance of the carrier network, or worse.

To configure GTP protocol anomalies, go to **Security Profiles > GTP Profile**, and edit a GTP profile. Expand the **Protocol Anomaly** option. See Protocol Anomaly prevention options.

The following are some examples:

- The GTP header specifies the length of the packet excluding the mandatory GTP header. In GTP version 0 (GSM 09.60), the mandatory GTP header size is 20 bytes, whereas GTP version 1 (GSM 29.060) specifies that the minimum length of the GTP header is 8 bytes. The GTP packet is composed of the header, followed by Information Elements typically presented in a Type-Length-Value format. It is possible for an attacker to create a GTP packet with a GTP header field length that is incompatible with the length of the necessary information elements.
- The same concepts are true for GTP version 2 headers even though there are different fields in them.
- It is similarly possible for an attacker to create a packet with an invalid IE length. Invalid lengths may cause protocol stacks to allocate incorrect amounts of memory, and thereby cause crashes or buffer overflows.

By default the FortiOS Carrier firewall detects these problems, as well as other protocol anomalies, and drops the packets. All protocol anomaly options are set to **Deny** by default. However, you can change the policy to allow them.

Configuring Anti-overbilling in FortiOS Carrier

GPRS over billing attacks can be prevented with a properly configured Carrier-enabled FortiGate unit.

Over billing can occur when a subscriber returns his IP address to the IP pool. Before the billing server closes it, the subscriber's session is still open and vulnerable. If an attacker takes control of the subscriber's IP address, he can send or receive data and the subscriber will be billed for the traffic.

Over billing can also occur when an available IP address is reassigned to a new mobile station (MS). Subsequent traffic by the previous MS may be forwarded to the new MS. The new MS would then be billed for traffic it did not initiate.

Anti-overbilling with FortiOS Carrier

The Carrier-enabled FortiGate unit can be configured to assist with anti-overbilling measures. These measures ensure that the customer is only billed for connection time and data transfer that they actually use.

Anti-overbilling on the Carrier-enabled FortiGate unit involves:

- the administrator configuring the over billing settings in the GTP profile to notify the Gi firewall when a GTP tunnel is deleted
- the unit clearing the sessions when the Gi firewall receives a notification from the Gn/Gp firewall about a GTP tunnel being deleted This way, the Gi firewall prevents over billing by blocking traffic initiated by other users.

The three locations to configure anti-overbilling options include:

- **Network > Interface** — Edit a specific interface. Towards the bottom of the **Edit Interface** page, in the **Status** section, you can toggle **Gi Gatekeeper**.
- **System > Settings** — In the **Gi Gatekeeper Settings** section, set the **Context ID** and **Port** that anti-overbilling will take place on.
- **Security Profiles > GTP Profile** — Edit a specific GTP Profile. In the **Anti-Overbilling** section, edit the **Gi Firewall IP address**, **Port**, **Interface** and **Security Context ID**, to use for anti-overbilling measures.

For detailed options, see [Anti-Overbilling options](#).

Logging events on the Carrier-enabled FortiGate unit

Logging on the Carrier-enabled FortiGate unit is just like logging on any other FortiOS unit. The only difference with FortiOS Carrier is that there are a few additional events that you can log beyond the regular ones. These additional events are covered here.

To change FortiOS Carrier specific logging event settings, go to **Security Profiles > GTP Profile** and edit a GTP profile. Expand the **Log** section to change the settings. For detailed options, see Log options.

The following information is contained in each log entry:

Timestamp	The time and date when the log entry was recorded
Source IP address	The sender's IP address.
Destination IP address	The receiver's IP address. The sender-receiver pair includes a mobile phone on the GPRS local network, and a device on a network external to the GPRS network, such as the Internet.
Tunnel Identifier (TID) Tunnel Endpoint Identifier (TEID)	An identifier for the start and endpoints of a GTP tunnel. This information uniquely defines all tunnels. It is important for billing information based on the length of time the tunnel was active and how much data passed over the tunnel.
Message type	For available message types, see Common message types on carrier networks.
Packet status	<p>What action was performed on the packet. This field matches the logging options while you are configuring GTP logging. See Logging events on the Carrier-enabled FortiGate unit on page 491.</p> <p>The status can be one of forwarded, prohibited, state-invalid, rate-limited, or tunnel-limited</p>
Virtual domain ID or name	A Carrier-enabled FortiGate unit can be divided into multiple virtual units, each being a complete and self-contained virtual FortiCarrier unit. This field indicates which virtual domain (VDOM) was responsible for the log entry. If VDOMs are not enabled on your unit, this field will be <code>root</code> .
Reason to be denied if applicable	If the packet that generated this log entry was denied or blocked, this field will include what part of FortiOS denied or blocked that packet. Such as firewall, antivirus, webfilter, or spamfilter.

An example of the above log message format is for a Tunnel deleted log entry. When a tunnel is deleted, the log entry contains the following information:

- Timestamp
- Interface name (if applicable)
- SGSN IP address (source IP)
- GGSN IP address (destination IP)
- Tunnel ID
- Tunnel duration time in seconds
- Number of messages sent to the SGSN
- Number of messages sent to the GGSN

GTP message type filtering

FortiOS Carrier supports message filtering in GTP by the type of message.

This section includes:

Common message types on carrier networks

Carrier networks include many types of messages — some concern the network itself, others are content moving across the network, and still others deal with handshaking, billing, or other administration based issues.

GTP contains two major parts GTP for the control plane (GTP-C) and GTP for user data tunnelling (GTP-U). Outside of those areas there are only unknown message types.

GTP-C messages

GTP-C contains the networking layer messages. These address routing, versioning, and other similar low level issues.

When a subscriber requests a Packet Data Protocol (PDP) context, the SGSN will send a create PDP context request GTP-C message to the GGSN giving details of the subscriber's request. The GGSN will then respond with a create PDP context response GTP-C message which will either give details of the PDP context actually activated or will indicate a failure and give a reason for that failure. This is a UDP message on port 212.

GTP-C message types include Path Management Messages, Location Management Messages, and Mobility Management Messages.

Path Management Messages

Path management is used by one GSN to detect if another GSN is alive, or if it has restarted after a failure.

The path management procedure checks if a given GSN is alive or has been restarted after a failure. In case of SGSN restart, all MM and PDP contexts are deleted in the SGSN, since the associated data is stored in a volatile memory. In the case of GGSN restart, all PDP contexts are deleted in the GGSN.

Tunnel Management Messages

The tunnel management procedures are used to create, update, and delete GTP tunnels in order to route IP PDUs between an MS and an external PDN via the GSNs.

The PDP context contains the subscriber's session information when the subscriber has an active session. When a mobile wants to use GPRS, it must first attach and then activate a PDP context. This allocates a PDP context data structure in the SGSN that the subscriber is currently visiting and the GGSN serving the subscriber's access point.

Tunnel management procedures are defined to create, update, and delete tunnels within the GPRS backbone network. A GTP tunnel is used to deliver packets between an SGSN and a GGSN. A GTP tunnel is identified in each GSN node by a TEID, an IP address, and a UDP port number.

Location Management Messages

The location-management procedure is performed during the network-requested PDP context activation procedure if the GGSN does not have an SS7 MAP interface (i.e., Gc interface). It is used to transfer location

messages between the GGSN and a GTP-MAP protocol-converting GSN in the GPRS backbone network.

Location management subprocedures are used between a GGSN that does not support an SS7 MAP interface (i.e., Gc interface) and a GTP-MAP protocol-conversing GSN. This GSN supports both Gn and Gc interfaces and is able to perform a protocol conversing between GTP and MAP.

Mobility Management Messages

The MM procedures are used by a new SGSN in order to retrieve the IMSI and the authentication information or MM and PDP context information in an old SGSN. They are performed during the GPRS attach and the inter-SGSN routing update procedures.

The MM procedures are used between SGSNs at the GPRS-attach and inter-SGSN routing update procedures. An identity procedure has been defined to retrieve the IMSI and the authentication information in an old SGSN. This procedure may be performed at the GPRS attach. A recovery procedure enables information related to MM and PDP contexts in an old SGSN to be retrieved. This procedure is started by a new SGSN during an inter-SGSN RA update procedure.

GTP-U messages

GTP-U is focused on user related issues including tunneling, and billing. GTP-U message types include MBMS messages, and GTP-U and Charging Management Messages

MBMS messages

Multimedia Broadcast and Multicast Services (MBMS) have recently begun to be offered over GSM and UMTS networks on UTRAN and GERAN radio access technologies. MBMS is mainly used for mobile TV, using up to four GSM timeslots for one MBMS connection. One MBMS packet flow is replicated by GGSN, SGSN and RNCs.

MBMS is split into the MBMS Bearer Service and the MBMS User Service. The MBMS User Service is basically the MBMS Service Layer and offers a Streaming- and a Download Delivery Method. The Streaming Delivery method can be used for continuous transmissions like Mobile TV services. The Download Method is intended for "Download and Play" services.

GTP-U and Charging Management Messages

SGSNs and GGSNs listen for GTP-U messages on UDP port 2152.

GTP' (GTP prime) is used for billing messages. It uses the common GTP messages (GTP Version Not Supported, Echo Request and Echo Response) and adds additional messages related to billing procedures.

Unknown Action messages

If the system doesn't know what type of message it is, it falls into this category. This is an important category of message because malformed messages may appear and need to be handled with security in mind.



Fortinet best practices dictate that you set **Unknown Action messages** to deny for security reasons.

Configuring message type filtering in FortiOS Carrier

GPRS Tunneling Protocol (GTP) is a group of IP-based communications protocols used to carry General Packet Radio Service (GPRS) traffic within Global System for Mobile Communications (GSM) and Universal Mobile Telecommunications System (UMTS) networks. It allows carriers to transport actual cellular packets over their network via tunneling.

In the CLI, there is a keyword for each type of GTP message for both message filtering, and for message rate limiting.



GTP message rate limiting is only accessible from the CLI using the command `configure firewall gtp`.

To configure GTP message type filtering - web-based manager

1. Go to **Security Profiles > GTP Profile**.
2. Select **Create New**.
3. Enter a name for this profile such as `msg_type_filtering`.
4. Select **Message Type Filtering** to expand it.
5. For each type of message in the list, select Allow or Deny. All messages are set to Allow by default.



Fortinet best practices dictate that the unknown message action should be set to **Deny** for security reasons as this will block malformed messages.

6. Optionally select and configure any other GTP features for this profile, such as logging.
7. Select **OK** to save the profile.
8. Apply the `msg_type_filtering` profile a security policy configured for GTP tunnel traffic.

To configure GTP message filtering and block Unknown Message Action messages- CLI

```
config firewall gtp
  edit msg_type_filtering
    config message-filter
      set unknown-message-action deny
    next
  end
end
```

Message Type Fields

Each of the following message types can be allowed or denied by your Carrier-enabled FortiGate unit depending on your carrier network and GTP traffic.

Unknown Message Action

Set this message type to deny.

Many attempts to hack into a carrier network will result in this unknown message type and therefore it is denied for security reasons.

Path Management Messages

Message Type	Used by	Description
Echo Request/Response	GTP-C, GTP-U, GTP'	Echo Request is sent on a path to another GSN to determine if the other node is alive. Echo Response is the reply.
Version not Supported	GTP-C, GTP-U, GTP'	There are multiple versions of GTP. Both devices communicating must use the same version of GTP, or this message will be the response.
Support Extension Headers Notification		Extensions are optional parts that a device can choose to support or not. If a device includes these extensions, it must include headers for the extensions to sure ensure proper formatting.

Tunnel Management Messages

Message Type	Used by	Description
Create PDP Context Request/ Response	GTP-C	Sent from an SGSN to a GGSN node as part of a GPRS PDP Context Activation procedure or the Network-Requested PDP Context Activation procedure. A valid request initiates the creation of a tunnel.
Update PDP Context Request/ Response	GTP-C	Used when PDP Context information changes, such as when a mobile device changes location.
Delete PDP Context Request/ Response	GTP-C	Used to terminate a PDP Context, and confirm the context has been deleted.
Create AA PDP Context Request/ Response	GTP-C	Sent as part of the GPRS Anonymous Access PDP Context Activation. It is used to create a tunnel between a context in the SGSN and a context in the GGSN.
Delete AA PDP Context Request/ Response	GTP-C	Sent as part of the GPRS PDP Anonymous Access Context Deactivation procedure to deactivate an activated PDP Context. It contains Cause and Private Extension Information Elements

Message Type	Used by	Description
Error Indication	GTP-U	<p>Sent to the GGSN when a tunnel PDU is received for the following conditions:</p> <ul style="list-style-type: none"> — No PDP context exists — PDP context is inactive — No MM context exists — GGSN deletes its PDP context when the message is received.
PDU Notification Request/ Response/ Reject Request/ Reject Response	GTP-C	<p>When receiving a Tunneled PDU (T-PDU), the GGSN checks if a PDP context is established for the given PDP address. If no PDP context has been established, the GGSN may initiate the Network-requested PDP Context Activation procedure by sending a PDU Notification Request to the SGSN.</p> <p>Reject Request - Sent when the PDP context requested by the GGSN cannot be established.</p>

Location Management Messages

Message Type	Used By	Description
Send Routing Information for GPRS Request/ Response	GTP-C	Sent by the GGSN to obtain location information for the MS. This message type contains the IMSI of the MS and Private Extension.
Failure Report Request/ Response	GTP-C	<p>Sent by the GGSN to the HLR when a PDU reject message is received.</p> <p>The GGSN requests the HLR to set the flag and add the GGSN to the list of nodes to report to when activity from the subscriber that owns the PDP address is detected.</p> <p>The message contains the subscriber IMSI and Private Extension</p>
Note MS GPRS Present Request/ Response	GTP-C	<p>When the HLR receives a message from a mobile with MDFG set, it clears the MDFG and sends the Note MS Present message to all GGSN's in the subscriber's list.</p> <p>This message type contains subscriber IMSI, GSN Address and Private Extension</p>

Mobility Management Messages

Message Type	Used By	Description
Identification Request/Response	GTP-C	Sent by the new SGSN to the old SGSN to request the IMSI for a MS when a GPRS Attach is done with a P-TMSI and the MS has changed SGSNs since the GPRS Detach was done.
SGSN context Request/ Response/ Acknowledge	GTP-C	Sent by the new SGSN to the old SGSN to request the MM and PDP Contexts for the MS.
Forward Relocation Request/ Response/ Complete/ Complete Acknowledge	GTP-C	<p>Indicates mobile activation/deactivation within a Routing Area. This prevents paging of a mobile that is not active (visited VLR rejects calls from the HLR or applies Call Forwarding). Note that the mobile station does not maintain an attach/detach state.</p> <p>SRNS contexts contain for each concerned RAB the sequence numbers of the GTP-PDUs next to be transmitted in uplink and downlink directions.</p>
Relocation Cancel Request/ Response	GTP-C	Send to cancel the relocation of a connection.
Forward SRNS Context/ Context Acknowledge	GTP-C	This procedure may be used to trigger the transfer of SRNS contexts from RNC to CN (PS domain) in case of inter system forward handover.
RAN Information Relay	GTP-C	<p>Forward the Routing Area Network (RAN) information.</p> <p>A Routing Area (RA) is a subset of a GSM Location Area (LA). A RA is served by only one SGSN. Ensures that regular radio contact is maintained by the mobile</p>

MBMS messages

Message Type	Used By	Description
MBMS Notification Request/ Response/ Reject Request/ Reject Response	GTP-C	Notification of the radio access devices.
Create MBMS Context Request/ Response	GTP-C	<p>Request to create an active MBMS context. The context will be pending until the response is received.</p> <p>Once active, the MBMS context allows the MS to receive data from a specific MBMS source</p>

Message Type	Used By	Description
Update MBMS Context Request/ Response	GTP-C	
Delete MBMS Context Request/ Response	GTP-C	Request to deactivate the MBMS context. When the response is received, the MBMS context will be inactive.

GTP-U and Charging Management Messages

Message Type	Used By	Description
G-PDU	GTP-C, GTP-U	GPRS Packet data unit delivery message.
Node Alive Request/Response	GTP-C, GTP-U	Used to inform rest of network when a node starts service.
Redirection Request/Response	GTP-C, GTP-U	Used to divert the flow of CDRs from the CDFs to another CGF when the sender is being removed, or they are used when the CGF has lost its connection to a downstream system.
Data Record Transfer Request/Response	GTP-C, GTP-U	Used to reliably transport CDRs from the point of generation (SGSN/GGSN) to non-volatile storage in the CGF

GTP identity filtering

FortiOS Carrier supports a number of filtering methods based on subscriber identity such as APN filtering, IMSI filtering, and advanced filtering.

This section includes:


IMSI on carrier networks

The International Mobile Subscriber Identity (IMSI) number is central to identifying users on a carrier network. It is a unique number that is assigned to a cell phone or mobile device to identify it on the GSM or UTM network.

Typical the IMSI number is stored on the SIM card of the mobile device and is sent to the network as required.

An IMSI number is 15 digits long, and includes the Mobile Country Code (MCC), Mobile Network Code (MNC), and Mobile Station Identification Number (MSIN).

IMSI codes



The Home Network Identity (HNI) is made up of the MCC and MNC. The HNI is used to fully identify a user's home network. This is important because some large countries have more than one country code for a single carrier. For example a customer with a mobile carrier on the East Coast of the United States would have a different MCC than a customer on the West Coast with the same carrier because even though the MNC would be the same the MCC would be different — the United States uses MCCs 310 to 316 due to its size.

If an IMSI number is not from the local carrier's network, IMSI analysis is performed to resolve the number into a Global Title which is used to access the user's information remotely on their home carrier's network for things like billing and international roaming.

Other identity and location based information elements

IMSI focuses on the user, their location, and carrier network. There are other numbers used to identify different user related Information Elements (IE).

These identity and location based elements include:

- Access Point Number (APN)
- Mobile Subscriber Integrated Services Digital Network (MSISDN)
- Radio Access Technology (RAT) type
- User Location Information (ULI)
- Routing Area Identifier (RAI)
- International Mobile Equipment Identity (IMEI)

Access Point Number (APN)

The Access Point Number (APN) is used in GPRS networks to identify an IP packet data network that a user wants to communicate with. The Network Identifier describes the network and optionally the service on that network that the GGSN is connected to. The APN also includes the MCC and MCN, which together locate the network the GGSN belongs to. An example of an APN in the Barbados using Digicel as the carrier that is connecting to the Internet is `internet.mcc342.mnc750.gprs`.

When you are configuring your Carrier-enabled FortiGate unit's GTP profiles, you must first configure the APN. It is critical to GTP communications and without it no traffic will flow.

The access point can then be used in a DNS query to a private DNS network. This process (called APN resolution) gives the IP address of the GGSN which serves the access point. At this point a PDP context can be activated.

Mobile Subscriber Integrated Services Digital Network (MSISDN)

This is a 15-digit number that, along with the IMSI, uniquely identifies a mobile user. Normally this number includes a 2-digit country code, a 3-digit national destination code, and a 10-digit subscriber number or the phone number of the mobile device, and because of that may change over time if the user changes their phone number. The MSISDN number follows the ITU-T E.164 numbering plan.

Radio Access Technology (RAT) type

The RAT type represents the radio technology used by the mobile device. This can be useful in determining what services or content can be sent to a specific mobile device. FortiOS Carrier supports:

- **UMTS Terrestrial Radio Access Network (UTRAN)**, commonly referred to as 3G, routes many types of traffic including IP traffic. This is one of the faster types.
- **GSM EDGE Radio Access Network (GERAN)** is a key part of the GSM network which routes both phone calls and data.
- **Wireless LAN (WLAN)** is used but not as widely as the other types. It is possible for the mobile device to move from one WLAN to another such as from an internal WLAN to a commercial hot spot.
- **Generic Access Network (GAN)** can also be called unlicensed mobile access (UMA). It routes voice, data, and SIP over IP networks. GAN is commonly used for mobile devices that have a dual-mode and can hand-off between GSM and WLANs.
- **High Speed Packet Access (HSPA)** includes two other protocols High Speed Downlink and Uplink Packet Access protocols (HSDPA and HSUPA respectively). It improves on the older WCDMA protocols by better using the radio bandwidth between the mobile device and the radio tower. This results in an increased data transfer rate for the user.

RAT type is part of advanced filtering configuration. See [Configuring advanced filtering in FortiOS Carrier](#).

User Location Information (ULI)

Gives Cell Global Identity/Service Area Identity (CGI/SAI) of where the mobile station is currently located. The ULI and the RAI are commonly used together to identify the location of the mobile device.

ULI is part of advanced filtering configuration. See [Configuring advanced filtering in FortiOS Carrier](#).

Routing Area Identifier (RAI)

Routing Areas (RAs) divide the carrier network and each has its own identifier (RAI). When a mobile device moves from one routing area to another, the connection is handled by a different part of the network. There are normally multiple cells in a routing area. There is only one SSGN per routing area. The RAI and ULI are commonly used to determine a user's location.

RAI is part of advanced filtering configuration. See [Configuring advanced filtering in FortiOS Carrier](#).

International Mobile Equipment Identity (IMEI)

IMEI is a unique 15-digit number used to identify mobile devices on mobile networks. It is very much like the MAC address of a TCP/IP network card for a computer. It can be used to prevent network access by a stolen phone — the carrier knows the mobile phone's IMEI, and when it is reported stolen that IMEI is blocked from accessing the carrier network no matter if it has the same SIM card as before or not. It is important to note that the IMEI stays with the mobile phone or device where the other information is either location based or stored on the removable SIM card.

IMEI type is part of advanced filtering configuration. See [Configuring advanced filtering in FortiOS Carrier](#).

When to use APN, IMSI, or advanced filtering

At first glance APN, IMSI, and advanced filtering have parts in common. For example two can filter on APN, and another two can filter on IMSI. The difficulty is knowing when to use which type of filtering.

Identity filtering comparison

Filtering type	Filter on the following data:	When to use this type of filtering
APN	APN	Filter based on GTP tunnel start or destination
IMSI	IMSI, MCC-MNC	Filter based on subscriber information
Advanced	PDP context, APN, IMSI, MSISDN, RAT type, ULI, RAI, IMEI	When you want to filter based on: <ul style="list-style-type: none"> • user phone number (MSISDN) • what wireless technology the user employed • to get on the network (RAT type) • user location (ULI and RAI) • handset ID, such as for stolen phones (IMEI)

APN filtering is very specific — the only identifying information that is used to filter is the APN itself. This will always be present in GTP tunnel traffic, so all GTP traffic can be filtered using this value.

IMSI filtering can use a combination of the APN and MCC-MNC numbers. The MCC and MNC are part of the APN, however filtering on MCC-MNC separately allows you to filter based on country and carrier instead of just the destination of the GTP Tunnel.

Advanced filtering can go into much deeper detail covering PDP contexts, MSISDN, IMEI, and more not to mention APN, and IMSI as well. If you can't find the information in APN or IMSI that you need to filter on, then use Advanced filtering.

Configuring APN filtering in FortiOS Carrier

To configure APN filtering go to **Security Profiles > GTP Profile**. Select a profile or create a new one, and expand **APN filtering**.



When you are configuring your Carrier-enabled FortiGate unit's GTP profiles, you must first configure the APN. It is critical to GTP communications and without it no traffic will flow.

Enable APN Filter	Select to enable filtering based on APN value.
Default APN Action	Select either Allow or Deny for all APNs that are not found in the list. The default is Allow.
Value	Displays the APN value for this entry. Partial matches are allowed using wildcard. For example <code>*.mcc333.mcn111.gprs</code> would match all APNs from country 333 and carrier 111 on the gprs network.
Mode	<p>Select one or more of the methods used to obtain APN values.</p> <p>Mobile Station provided - The APN comes from the mobile station where the mobile device connected. This is the point of entry into the carrier network for the user's connection.</p> <p>Network provided - The APN comes from the carrier network.</p> <p>Subscription Verified - The user's subscription has been verified for this APN. This is the most secure option.</p>
Action	One of allow or deny to allow or block traffic associated with this APN.
Delete icon	Select to remove this APN entry from the list.
Edit icon	Select to change the information for this APN entry.
Add APN	<p>Select to add an APN to the list. Not active while creating GTP profile, only when editing an existing GTP profile.</p> <p>Save all changes before adding APNs. A warning to this effect will be displayed when you select the Add APN button.</p>

The Add APN button is not activated until you save the new GTP profile. When you edit that GTP profile, you will be able to add new APNs.

Configuring IMSI filtering in FortiOS Carrier

In many ways the IMSI on a GPRS network is similar to an IP address on a TCP/IP network. Different parts of the number provide different pieces of information. This concept is used in IMSI filtering on FortiOS Carrier.

To configure IMSI filtering go to **Security Profiles > GTP Profile** and expand **IMSI filtering**.

While both the APN and MCC-MCN fields are optional, without using one of these fields the IMSI entry will not be useful as there is no information for the filter to match.

Enable IMSI Filter	Select to turn on IMSI filtering.
Default IMSI Action	<p>Select Allow or Deny. This action will be applied to all IMSI numbers except as indicated in the IMSI list that is displayed.</p> <p>The default value is Allow.</p>
APN	<p>The Access Point Number (APN) to filter on.</p> <p>This field is optional.</p>
MCC-MNC	<p>The Mobile Country Code (MCC) and Mobile Network Code (MNC) to filter on. Together these numbers uniquely identify the carrier and network of the GGSN being used.</p> <p>This field is optional.</p>
Mode	<p>Select the source of the IMSI information as one or more of the following:</p> <p>Mobile Station provided - the IMSI number comes from the mobile station the mobile device is connecting to.</p> <p>Network provided - the IMSI number comes from the GPRS network which could be a number of sources such as the SGSN, or HLR.</p> <p>Subscription Verified - the IMSI number comes from the user's home network which has verified the information.</p> <p>While Subscription Verified is the most secure option, it may not always be available. Selecting all three options will ensure the most complete coverage.</p>
Action	Select the action to take when this IMSI information is encountered. Select one of Allow or Deny.
Delete Icon	Select the delete icon to remove this IMSI entry.
Edit Icon	Select the edit icon to change information for this IMSI entry.
Add IMSI	<p>Select to add an IMSI to the list. Not active while creating GTP profile, only when editing an existing GTP profile.</p> <p>Save all changes before adding IMSIs. A warning to this effect will be displayed when you select the Add IMSI button.</p>

Configuring advanced filtering in FortiOS Carrier

Compared to ADN or IMSI filtering, advanced filtering is well named. Advanced filtering can be viewed as a catch-all filtering option — if ADN or IMSI filtering doesn't do what you want, then advanced filtering will. The advanced filtering can use more information elements to provide considerably more granularity for your filtering.

Enable	Select to turn on advanced filtering.
Default Action	Select Allow or Deny as the default action to take when traffic does not match an entry in the advanced filter list .
Messages	<p>Optionally select one or more types of messages this filter applies to:</p> <p>Create PDP Context Request, Create PDP Context Response, Update PDP Context Request, or Update PDP Context Response.</p> <p>Selecting Create PDP Context Response or Update PDP Context Response limits RAT type to only GAN and HSPA, and disables the APN, APN Mode, IMSI, MSISDN, ULI, RAI, and IMEI fields.</p> <p>To select Update PDP Context Request, APN Restriction must be set to all. Selecting Update PDP Context Request disables the APN, MSISDN, and IMEI fields.</p> <p>if all message types are selected, only the RAT Types of GAN and HSPA are available to select.</p>
APN Restriction	APN Restriction either allows all APNs or restricts the APNs to one of four categories — Public-1, Public-2, Private-1, or Private-2. This can also be combined with a specific APN or partial APN as well as specifying the APN mode.
RAT Type	Select one or more of the Radio Access Technology Types listed. These fields control how a user accesses the carrier's network. You can select one or more of UTRAN, GERAN, WLAN, GAN, HSPA, or any.
ULI	<p>The user location identifier. Often the ULI is used with the RAI to locate a user geographically on the carrier's network.</p> <p>The ULI is disabled when Create PDP Context Response or Update PDP Context Response messages are selected.</p>
RAI	<p>The router area identifier. There is only one SGSN per routing area on a carrier network. This is often used with ULI to locate a user geographically on a carrier network.</p> <p>The RAI is disabled when Create PDP Context Response or Update PDP Context Response messages are selected.</p>

IMEI	<p>The International Mobile Equipment Identity. The IMEI uniquely identifies mobile hardware, and can be used to block stolen equipment.</p> <p>The IMEI is only available when Create PDP Context Request or no messages are selected.</p>
Action	<p>Select Allow or Deny as the action when this filter matches traffic.</p> <p>The default is Allow.</p>
Delete Icon	Select to delete this entry from the list.
Edit Icon	Select to edit this entry.
Add	<p>Select to add an advanced filter to the list. Not active while creating GTP profile, only when editing an existing GTP profile.</p> <p>Save all changes before adding advanced filters. A warning to this effect will be displayed when you select the Add button.</p>

SCTP Concepts

As of FortiOS version 5.0, the FortiGate natively handles SCTP (Stream Control Transport Protocol) traffic, as an alternative to TCP and UDP for use in Carrier networks. The FortiGate handles SCTP as if it would any other traffic.

Overview

SCTP is a connection-oriented transport protocol that overcomes some of the limitations of both TCP and UDP that prevent reliable transfer of data over IP-based networks (such as those used by telephony systems and carrier networks). The 'Stream' in SCTP refers to the sequence of user messages or packets that are considered at the same time to be individual objects and also treated as a whole by networked systems. SCTP is less vulnerable to congestion and flooding due to more advanced error handling and flood protection built into the protocol.

SCTP features as compared to TCP and UDP

Feature	SCTP	TCP	UDP
State required at each endpoint	yes	yes	no
Reliable data transfer	yes	yes	no
Congestion control and avoidance	yes	yes	no
Message boundary conservation	yes	no	yes

Feature	SCTP	TCP	UDP
Path MTU discovery and message fragmentation	yes	yes	no
Message bundling	yes	yes	no
Multi-homed hosts support	yes	no	no
Multi-stream support	yes	no	no
Unordered data delivery	yes	no	yes
Security cookie against SYN flood attack	yes	no	no
Built-in heartbeat (reachability check)	yes	no	N/A

All of these features are built into the design of the Protocol, and the structure of SCTP packets and networks. The FortiGate unit interprets the traffic and provides the necessary support for maintenance and verification features, but the features are not FortiGate specific. These features are documented in greater detail below.

State required at each endpoint

Constant back and forth acknowledgement and content verification messages are sent between all SCTP peer endpoints, and all endpoints' state machine actions must be synchronized for traffic to flow.

Reliable data transfer

SCTP places data and control information (eg. source, destination, verification) into separate messages, both sharing the same header in the same SCTP packet. This allows for constant verification of the contained data at both ends and along the path, preventing data loss or fragmentation. As well, data is not sent in an interruptible stream as in TCP.

Congestion control and avoidance

Built-in, constantly updating path detection and monitoring automatically redirect packets along alternate paths in case of traffic congestion or inaccessible destinations. For deliberate/malicious congestion control, see the below section on [Security cookie against SYN flood attack](#).

Message boundary conservation

SCTP is designed in such a way that no matter how messages are divided, redirected, or fragmented, the message boundaries will be maintained within the packets, and all messages cannot be appended without tripping verification mechanisms.

Path MTU discovery and message fragmentation

SCTP is capable of Path Maximum Transmission Unit discovery, as outlined in RFC4821. Two specific alterations have been made to how SCTP handles MTU. First, that endpoints will have separate MTU estimates for each possible multi-homed endpoint. Second, that bundled message fragments (as explained below) will be directed based on MTU calculations, so that retransmissions (if necessary) will be sent without delay to alternate addresses.

Message bundling

SCTP is a message-oriented protocol, which means that despite being a streaming data protocol, it transports a sequence of specific messages, rather than transporting a stream of bytes (like TCP). Since some data transmissions are small enough to not require a complete message's worth of content, so multiple pieces of content will be transmitted simultaneously within the messages.

Multi-homed hosts support

SCTP supports multi-homing, which is a network structure in which one or multiple sources/destinations has more than one IP address. SCTP can adapt to multi-homing scenarios and redirect traffic to alternate IP addresses in case of failure.

Multi-stream support

Due to the message bundling feature allowing for multiple pieces of content to be sent in messages at once, SCTP can 'multi-stream' content, by deliberately dividing it among messages at a fixed rate, so that multiple types of content (eg. both images and text) can be loaded at once, at the same pace.

Unordered data delivery

With control messages in every packet to provide verification of any packet's data and its place in the stream, the data being transmitted can actually arrive in any order, and verify that all has arrived or that some is missing.

Security cookie against SYN flood attack

Since every packet contains verification of its place in the stream, it makes it easy for the protocol to detect when redundant, corrupted or malicious packets flood the path, and they are automatically dropped when necessary.

Built-in heartbeat (reachability check)

Endpoints automatically send specific control chunks among the other SCTP packet information to peer endpoints, to determine the reachability of the destination. Heartbeat acknowledgement packets are returned if the destination is available.

SCTP Firewall

FortiGate stateful firewalls will protect and inspect SCTP traffic, according to RFC4960. SCTP over IPsec VPN is also supported. The FortiGate device is inserted as a router between SCTP endpoints. It checks SCTP Syntax for the following information:

- Source and destination port
- Verification Tag
- Chunk type, chunk flags, chunk length
- Sequence of chunk types
- Associations

The firewall also oversees and maintains several SCTP security mechanisms:

- SCTP four-way handshake
- SCTP heartbeat
- NAT over SCTP

The firewall has IPS DoS protection against known threats to SCTP traffic, including INIT/ACK flood attacks, and SCTP fuzzing.

Troubleshooting

This section offers troubleshooting options for Carrier-related issues.

This section includes:

FortiOS Carrier diagnose commands

This section includes diagnose commands specific to FortiOS Carrier features such as GTP.

GTP related diagnose commands

This CLI command allows you to gain information on GTP packets, logs, statistics, and other information.

```
diag firewall gtp <command>
```

apn list <gtp_profile>	The APN list entries in the specified GTP profile
auth-ggsns show <gtp_profile>	The authorized GGSNs entries for the specified GTP profile. Any GGSNs not on this list will not be recognized.
auth-sgsns show <gtp_profile>	The authorized SGSNs list entries for the specified GTP profile. Any SGSNs not on this list will not be recognized.
handover-grp show <gtp_profile>	The handover group showing the range of allowed handover group IP addresses. The handover group acts like a white list of allowed GTP addresses with a default deny at the end — if the GTP address is not on the list, it is denied.
ie-remove-policy list <gtp_profile>	List of IE policies in the IE removal policy for this GTP profile. The information displayed includes the message count for this policy, the length of the SGSN, the list of IEs, and list of SGSN IP addresses.
imsi list <gtp_profile>	IMSI filter entries for this GTP profile. The information displayed includes the message count for this filter, length of the IMSI, the length of the APN and IMSI, and of course the IMSI and APN values.
invalid-sgsns-to-long list <gtp_profile>	List of SGSNs that do not match the filter criteria. These SGSNs will be logged.
ip-policy list <gtp_profile>	List the IP policies including message count for each policy, the action to take, the source and destination IP addresses or ranges, and masks.
noip-policy <gtp_profile>	List the non-IP policies including the message count, which mode, the action to take, and the start and end protocols to be used by decimal number.

	Select list or flush.
path {list flush}	List the GTP related paths in FortiOS Carrier memory. Flush the GTP related paths from memory.
policy list <gtp_policy>	The GTP advanced filter policy information for this GTP profile. The information displayed for each entry includes a count for messages matching this filter, a hexadecimal mask of which message types to match, the associated flags, action to take on a match, APN selection mode, MSISDN, RAT types, RAI, ULI, and IMEI.
profile list	Displays information about the configured GTP profiles. You will not be able to see the bulk of the information if you do not log the output to a file.
runtime-stat flush	Select to flush the GTP runtime statistics from memory.
stat	Display the GTP runtime statistics — details on current GTP activity. This information includes how many tunnels are active, how many GTP profiles exist, how many IMSI filter entries, how many APN filter entries, advanced policy filter entries, IE remove policy filter entries, IP policy filter entries, clashes, and dropped packets.
tunnel {list flush}	Select one of list or flush. List lists all the GTP tunnels currently active. Flush clears the list of active GTP tunnels. This does not clear the clash counter displayed in the <code>stat</code> command.

Applying IPS signatures to IP packets within GTP-U tunnels

GTP-U (GTP user data tunnelling) tunnels carry user data packets, signaling messages and error information. GTP-U uses UDP port 2152. Carrier-enabled FortiGate units can apply IPS intrusion protection and detection to GTP-U user data sessions.

To apply IPS to GTP-U user data sessions, add an IPS Sensor to a profile and add the profile to a security policy that accepts GTP-U tunnels. The security policy Service field must be set to GTP or ANY to accept GTP-U packets.

The Carrier-enabled FortiGate unit intercepts packets with destination port 2152, removes the GTP header and handles the packets as regular IP packets. Applying an IPS sensor to the IP packets, the Carrier-enabled FortiGate unit can log attacks and pass or drop packets depending on the configuration of the sensor.

If the packet is GTP-in-GTP, or a nested tunnel, the packets are passed or blocked without being inspected.

To apply an IPS sensor to GTP-U tunnels

1. Go to **Security Profiles > Intrusion Protection** and select Create New (+) to add an IPS Sensor.
2. Configure the IPS Sensor to detect attacks and log, drop, or pass attack packets.

See the Intrusion Protection section of the [FortiOS UTM Guide](#).

3. Go to **Policy & Objects > IPv4 Policy** and apply the IPS sensor to the security policy.
4. Go to **Policy & Objects > IPv4 Policy** and select Create New to add a security policy or select a security policy.
5. Configure the security policy to accept GTP traffic.
In the security policy configure the source and destination settings to match the GTP traffic. Service to GTP or ANY so that the security policy accepts GTP traffic.
6. Select the GTP profile within the security policy.
7. Configure any other required security policy settings.
8. Select **OK** to save the security policy.

GTP packets are not moving along your network

When GTP packets are not getting to their destination, this could be caused by any one of a number of issues. General troubleshooting principals apply here.

The following sections provide some suggestions on how to troubleshoot this issue:

- [Attempt to identify the section of your network with the problem](#)
- [Ensure you have an APN configured](#)
- [Check the logs and adjust their settings if required](#)
- [Check the routing table](#)
- [Perform a sniffer trace](#)
- [Generate specific packets to test the network](#)

Attempt to identify the section of your network with the problem

The first step is to determine how widespread this problem is. Does it affect the whole GPRS network, or just one or two devices?

If the entire network is has this problem, the solution is likely a more general one such as ensuring the security policies allow GTP traffic to pass, the GTP profile specifies SSGNs and GSGNs, or ensuring the GTP general settings are not overly limiting.

If one part of the network is affected, the problem is more likely centered around configurations with those network devices specified such as the handover group, or authorized SGSNs/GGSNs. It is also possible that small portions of the network may have hardware related issues such as cabling or faulty hardware. This section does not address those issues, and assumes hardware is not the problem.

The handover group is a white list of GTP addresses allowed to handle GTP messages. If a device's address is not on this list, it will be denied.

Ensure you have an APN configured

When you configure your GTP profile, ensure you first configure the APN. Without it, there will be no flow of traffic. The APN is used in nearly all GTP communications and without it, the Carrier-enabled FortiGate unit doesn't have the information it needs.

Check the logs and adjust their settings if required

During normal operation, the log settings will show any problems on the network but may not provide the level of details required to fully troubleshoot the problem. The reason for this is that the level of detail required for troubleshooting would quickly overwhelm the daily logs without any real benefit.

GTP related events in the event log will have message IDs in the range 41216 to 41222. For more information on GTP log messages, see the Log Message Reference. For more information on logging in general, see the Logging and Reporting guide.

Once there is a problem to troubleshoot, check the logs to trace the traffic patterns and narrow down the possible sources of the problem. There may be enough detail for you to locate and fix the problem without changing the log settings.



Remember to set any changes you made to the log settings back to their original values when you are done troubleshooting. Otherwise, the amount of detail will overwhelm your logging.

However, if more detail is required you can change settings such as:

- Lower the Log Frequency number in GTP Profiles so fewer or no log messages are dropped. This will allow a more accurate picture of everything happening on the network, where you may have had only a partial picture before.
- Ensure all the GTP log events are enabled to provide you with a complete picture.
- Ensure that all relevant event types are enabled under **Log & Report > Log Config > Log Settings**.

For more information on GTP related logging, see Logging events on the Carrier-enabled FortiGate unit.

General information to look for in the logs includes:

- Are all packets having problems or just certain types?
- Are all devices on the network having problem, or just certain devices?
- Is it just GTP traffic that is having problems or are all types of traffic having the same problem?

Check the routing table

On any network, the routing table determines how packets reach their destination. This is also true on a carrier network.

If the Carrier-enabled FortiGate unit is running in NAT mode, verify that all desired routes are in the routing table — local subnets, default routes, specific static routes, and dynamic routing protocols. For complete information, it is best to check the routing table in the CLI. This method provides more complete information.



If VDOMs are enabled on your Carrier-enabled FortiGate unit, all routing related CLI commands must be performed within a VDOM and not in the global context.

To check the routing table using the CLI

```
# get router info routing-table all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
```



```
* - candidate default

S* 0.0.0.0/0 [10/0] via 192.168.183.254, port2
S 1.0.0.0/8 [10/0] via 192.168.183.254, port2
S 2.0.0.0/8 [10/0] via 192.168.183.254, port2
C 10.142.0.0/23 is directly connected, port3
B 10.160.0.0/23 [20/0] via 10.142.0.74, port3, 2d18h02m
C 192.168.182.0/23 is directly connected, port2
```

Examining an entry from the routing table above:

```
B 10.160.0.0/23 [20/0] via 10.142.0.74, port3, 2d18h02m
```

B	BGP. The routing protocol used.
10.160.0.0/23	The destination of this route including netmask.
[20/0]	20 indicates administrative distance of 20 out of a range of 0 to 255. 0 is an additional metric associated with this route, such as in OSPF
10.142.0.74	The gateway, or next hop.
port3	The interface used by this route.
2d18h02m	How old this route is, in this case almost three days old.

Perform a sniffer trace

When troubleshooting network traffic, it helps to look inside the headers of packets to determine if they are traveling along the route you expect. Packet sniffing can also be called a network tap, packet capture, or logic analyzing.



If your Carrier-enabled FortiGate unit has NP interfaces that are offloading traffic, this will change the sniffer trace. Before performing a trace on any NP interfaces, disable offloading on those interfaces.

What can sniffing packets tell you

If you are running a constant traffic application such as ping, packet sniffing can tell you if the traffic is reaching the destination, what the port of entry is on the Carrier-enabled FortiGate unit, if the ARP resolution is correct, and if the traffic is being sent back to the source as expected.

Sniffing packets can also tell you if the Carrier-enabled FortiGate unit is silently dropping packets for reasons such as RPF (Reverse Path Forwarding), also called Anti Spoofing. This prevents an IP packet from being forwarded if its source IP address either does not belong to a locally attached subnet (local interface), or be a hop on the routing between the FortiOS Carrier and another source (static route, RIP, OSPF, BGP). Note that RPF can be disabled by turning on asymmetric routing in the CLI (`config system setting, set asymmetric enable`), however this will disable stateful inspection on the Carrier-enabled FortiGate unit and consequently cause many features to be turned off.



If you configure virtual IP addresses on your Carrier-enabled FortiGate unit, the unit will use those addresses in preference to the physical IP addresses. If not configured properly, secondary IP addresses can cause a broadcast storm. You will notice the secondary address being preferred when you are sniffing packets because all the traffic will be using the virtual IP addresses. This is due to the ARP update that is sent out when the VIP address is configured.

How to sniff packets

The general form of the internal FortiOS packet sniffer command is:

```
diag sniffer packet <interface_name> <'filter'> <verbose> <count>
```

To stop the sniffer, type `CTRL+C`.

<interface_name>	The name of the interface to sniff, such as <code>port1</code> or <code>internal</code> . This can also be <code>any</code> to sniff all interfaces.
<'filter'>	What to look for in the information the sniffer reads. <code>none</code> indicates no filtering, and all packets will be displayed as the other arguments indicate. The filter must be inside single quotes (').
<verbose>	The level of verbosity as one of: <ul style="list-style-type: none"> 1 - print header of packets 2 - print header and data from IP of packets 3 - print header and data from Ethernet of packets
<count>	The number of packets the sniffer reads before stopping. If you don't put a number here, the sniffer will run forever until you stop it with <code><CTRL C></code> .

For a simple sniffing example, enter the CLI command `diag sniffer packet port1 none 1 3`. This will display the next 3 packets on the `port1` interface using no filtering, and using verbose level 1. At this verbosity level you can see the source IP and port, the destination IP and port, action (such as `ack`), and sequence numbers.

In the output below, port 443 indicates these are HTTPS packets, and 172.20.120.17 is both sending and receiving traffic.

```
Head_Office_620b # diag sniffer packet port1 none 1 3
interfaces=[port1]
filters=[none]
0.545306 172.20.120.17.52989 -> 172.20.120.141.443: psh 3177924955 ack 1854307757

0.545963 172.20.120.141.443 -> 172.20.120.17.52989: psh 1854307757 ack 3177925808

0.562409 172.20.120.17.52988 -> 172.20.120.141.443: psh 4225311614 ack 3314279933
```

Generate specific packets to test the network

If some packets are being delivered as expected while others are not, or after you believe you have fixed the problem, it is a good idea to generate specific traffic to test your network.

For example if you discover through log messages and packet sniffing that Create PDP Context Request messages are not being delivered between two SGSNs, you can generate those specific messages on your network to confirm they are the problem, and later that you have solved the problem and they are now being delivered as expected.

This step requires a third party traffic generation tool, either hardware or software. This is not supported by Fortinet.

Chapter 6 - Firewall

This guide contains a number of different topics that, at its simplest, can be grouped into fundamental firewall topics such as policies, objects and network defense and topics that have to do with the optimization of the firewall such as WAN optimization, proxies and caching.

Fundamentals

"[Firewall concepts](#)" on [page 522](#) explains the ideas behind the components, techniques and processes that are involved in setting up and running a firewall in general and the FortiGate firewall in particular. The premise here is that regardless of how experienced someone is with firewalls as they go through the process of configuring a firewall that is new to them they are likely to come across a term or setting that they may not be familiar with even if it is only in the context of the setting they are working in at the moment. FortiGate firewalls are quite comprehensive and can be very granular in the functions that they perform, so it makes sense to have a consistent frame of reference for the ideas that we will be working with.

Some examples of the concepts that will be addressed here are:

- "What is a Firewall?"
- "NAT"
- "IPv6"

"Firewall objects" describes the following firewall objects:

- Addressing
- Services
- Firewall Policies

"[Network defense](#)" on [page 665](#) describes various methods of defending your Network using the abilities of the FortiGate Firewall.

"[Object configuration](#)" on [page 705](#) is similar to a cookbook in that it will refer to a number of common tasks that you will likely perform to get the full functionality out of your FortiGate firewall. Because of the way that firewalls are designed, performing many of the tasks requires that firewall components be set up in a number of different sections of the interface and be configured to work together to achieve the desired result. This section will bring those components all together as a straight forward series of instructions.

FortiGate firewall components

The FortiGate firewall is made up of a number of different components that are used to build an impressive list of features that have flexibility of scope and granularity of control that provide protection that is beyond that provided by the basic firewalls of the past.

Some of the components that FortiOS uses to build features are:

- Interfaces
- VLANs
- Soft Switches

- Zones
- Predefined Addresses
- IP address based
- FQDN based
- Geography based
- Access Schedules
- Authentication
- Local User based
- Authentication Server based (Active Directory, Radius, LDAP)
- Device Based
- Configureable Services
- IPv4 and IPv6 protocol support

The features of FortiOS include but are not limited to:

- Security profiles, sometimes referred to as Unified Threat Management (UTM) or Next Generation Firewall (NGFW)
- Predefined firewall addresses (this includes IPv4 and IPv6, IP pools, . wildcard addresses and netmasks, and geography-based addresses)
- Monitoring traffic
- Traffic shaping and per-IP traffic shaping (advanced)
- Firewall schedules
- Services (such as AOL, DHCP and FTP)
- Logging traffic
- Quality of Service (QoS)
- Identity-based policies
- Endpoint security

Firewall optimization

There are a few different methodologies of optimization and most of these methodologies has been divided into:

- Concepts section - This will have the basic ideas behind the how and why of the topic. Because the number of topics is larger, the ideas are not as pervasive and the content is not so extensive as in the Fundamental section, some of the topics will include instructions on the configuration of that individual topic in order to keep the information for granular topics together.
- Configuration section- Just like the Configuration section of the Fundamentals, this will be a cookbook style of documentation showing how to configure something that achieves a specific functionality from the FortiGate.

The optimization topics include:

- [Secure web gateway, WAN optimization, web caching and WCCP](#)
- [Example topologies relevant to WAN optimization](#)
- [Inside FortiOS: WAN optimization](#)
- [WAN optimization concepts](#)
- [WAN optimization configuration](#)

- Peers and authentication groups
- Web cache concepts
- Web cache configuration
- WCCP concepts
- WCCP configuration
- Web proxy concepts
- Web proxy configuration
- Explicit proxy concepts
- Explicit proxy configuration
- Transparent proxy concepts
- Transparent proxy configuration
- FTP proxy concepts
- FTP proxy configuration
- Diagnose commands for WAN optimization

How does a FortiGate protect your network?

The FortiGate firewall protects your network by taking the various components and using them together to build a kind of wall or access control point so that anyone that is not supposed to be on your network is prevented from accessing your network in anyway other than those approved by you. It also protects your network from itself by keeping things that shouldn't happen from happening and optimizing the flow of traffic so that the network is protected from traffic congestion that would otherwise impede traffic flow.

Most people have at one time or another played with the children's toy system that is made up of interlocking blocks. The blocks come in different shapes and sizes so that you can build structures to suit your needs and in your way. The components of the FortiGate firewall are similar. You are not forced to use all of the blocks all of the time. You mix and match them to get the results that you are looking for. You can build a very basic structure that's only function is to direct traffic in and out to the correct subnets or you can build a fortress that only allows specific traffic to specific hosts from specific hosts at specific times of day and that is only if they provide the credentials that have been pre-approved and all of the traffic is encrypted so that even when the traffic is out on the Internet it is private from the world. Just like the interlocking blocks, what you build is up to you, but chances are if you put them together the right way there isn't much that can't be built.

Here is one example of how the components could be put together to support the requirements of a network infrastructure design.

- Off the Internal interface you could have separate VLANs. One for each for the departments of Sales, Marketing and Engineering so that the traffic from the users on one VLAN does not intrude upon the hosts of the other VLANs and the department are isolated from one another for security reasons.
- To ease in the administration each of the VLAN sub-interfaces is made a member of a zone so that security policies that apply to all of the hosts on all of the VLANs can be applied to all of them at once.
- Using the addresses component each of the IP address ranges could be assigned a user friendly name so that they could be referred to individually and then for policies that would refer to them all as a whole the individual ranges to be made members of an address group.
- Firewall schedules could be created to address the differing needs of each of the groups so that Sales and Marketing could be allowed access to the Internet during regular business hours and the Engineering department could be allowed access during the lunch break.

- By setting up the outgoing policies to use FortiGuard Web-filtering the employees could be prevented from visiting inappropriate sites and thus enforcing the policies of the HR department.
- A couple of virtual IP addresses with port forwarding could be configured to allow users on the Internet to access a web server on the DMZ subnet using the company's only Public IP address without affecting the traffic that goes to the company's mail server that is hosted on a complete different computer.
- Even though the Web server on the same DMZ has an FTP service to allow for the uploading of web pages to the web server from the Marketing and Engineer teams, by placing a DENY policy on any FTP traffic from the Internet malicious users are prevented from abusing the FTP service.
- By monitoring the traffic as it goes through the policies you can verify that the policies are in working order.
- By using a combination of ALLOW and DENY policies and placing them in the correct order you could arrange for an outside contractor to be allowed to update the web site as well

These set of configurations is not extensive but it does give an idea of how different components can be mixed and matched to build a configuration that meets an organization's needs but at the same time protect it from security risks.

What's new for Firewall in 6.0.1

The following list contains new firewall features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- [IPv6 Neighbor Discovery Proxy on page 636](#)

What's new for Firewall in 6.0

The following list contains new firewall features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- [SSH MITM deep inspection on page 554](#)
- [Creating NAT46 IP pool and multiple \(secondary\) NAT64 prefixes on page 741](#)
- [Application groups for NGFW policies on page 532](#)
- [Wildcard FQDNs for SSL deep inspection exemptions on page 716](#)
- [IPv6 FQDN firewall addresses on page 718](#)
- [Firewall IPv6 address templates on page 719](#)
- [Port block allocation timeout on page 739](#)
- [WAN Optimization and web cache improvements. Changes can be found at:](#)
 - [WANopt storage on page 785](#)
 - [WANopt cache service on page 786](#)
 - [Video caching on page 787](#)
 - [diagnose wad csvc on page 915](#)
 - [diagnose wad worker on page 916](#)
- [SSL mirroring for policies on page 704](#)
- [ISDB and IRDB in firewall policies on page 686](#)
- [IPv6 support for GRE tunnels on page 631](#)

- [Dispatching traffic to WAD worker based on source affinity on page 756](#)
- [Explicit proxy authentication timeout on page 1](#)
- [Multiple web proxy PAC files in one VDOM on page 862](#)
- [Logging options in web proxy profiles on page 859](#)
- [Encryption strength for proxied SSH sessions on page 558](#)
- [DNS traffic in NGFW policy-mode on page 610](#)
- [Using a FortiCache as a cache service on page 834](#)
- [IPv6 SSH on page 640](#)
- [WCCP Cisco to FortiGate client using L2-forwarding tunneling on page 835](#)

Firewall concepts

There are a number of foundational concepts that are necessary to have a grasp of before delving into the details of how the FortiGate firewall works. Some of these concepts are consistent throughout the firewall industry and some of them are specific to more advanced firewalls such as the FortiGate. Having a solid grasp of these ideas and terms can give you a better idea of what your FortiGate firewall is capable of and how it will be able to fit within your networks architecture.

This chapter describes the following firewall concepts:

- [What is a firewall?](#)
- [FortiGate modes](#)
- [How packets are handled by FortiOS](#)
- [Interfaces and zones](#)
- [Access control lists](#)
- [IPv6](#)
- [NAT](#)
- ["IP pools" on page 572](#)
- [Services and TCP ports](#)
- ["Firewall policies" on page 529](#)
- ["SSL/SSH inspection" on page 553](#)
- ["VPN policies" on page 601](#)
- ["DSRI" on page 601](#)
- ["Interface policies" on page 602](#)
- ["Local-In policies" on page 608](#)
- ["Security policy 0" on page 609](#)
- ["Deny policies" on page 610](#)
- ["Accept policies" on page 610](#)
- [IPv6 Policies](#)
- ["Fixed port" on page 610](#)
- ["Endpoint security" on page 611](#)
- ["Traffic logging" on page 611](#)

What is a firewall?

The term firewall originally referred to a wall intended to confine a fire or potential fire within a building. Later uses refer to similar structures, such as the metal sheet separating the engine compartment of a vehicle or aircraft from the passenger compartment.

A firewall can either be software-based or hardware-based and is used to help keep a network secure. Its primary objective is to control the incoming and outgoing network traffic by analyzing the data packets and determining whether it should be allowed through or not, based on a predetermined rule set. A network's firewall builds a bridge between an internal network that is assumed to be secure and trusted, and another network, usually an external (inter)network, such as the Internet, that is not assumed to be secure and trusted.

Network layer or packet filter firewalls

Stateless firewalls

Stateless firewalls are the oldest form of these firewalls. They are faster and simple in design requiring less memory because they process each packet individually and don't require the resources necessary to hold onto packets like stateful firewalls. Stateful firewalls inspect each packet individually and check to see if it matches a predetermined set of rules. According to the matching rule the packet is either be allowed, dropped or rejected. In the case of a rejection an error message is sent to the source of the traffic. Each packet is inspected in isolation and information is only gathered from the packet itself. Simply put, if the packets were not specifically allowed according to the list of rules held by the firewall they were not getting through.

Stateful Firewalls

Stateful firewalls retain packets in memory so that they can maintain context about active sessions and make judgments about the state of an incoming packet's connection. This enables Stateful firewalls to determine if a packet is the start of a new connection, a part of an existing connection, or not part of any connection. If a packet is part of an existing connection based on comparison with the firewall's state table, it will be allowed to pass without further processing. If a packet does not match an existing connection, it will be evaluated according to the rules set for new connections. Predetermined rules are used in the same way as a stateless firewall but they can now work with the additional criteria of the state of the connection to the firewall.

Best Practices Tip for improving performance:



Blocking the packets in a denied session can take more cpu processing resources than passing the traffic through. By putting denied sessions in the session table, they can be kept track of in the same way that allowed session are so that the FortiGate unit does not have to redetermine whether or not to deny all of the packets of a session individually. If the session is denied all packets of that session are also denied.

In order to configure this you will need to use 2 CLI commands

```
config system setting
    set ses-denied-traffic enable
    set block-session-timer <integer 1 - 300> (this determines in
seconds how long, in seconds, the session is kept in the table)
end
```

Application layer firewalls

Application layer filtering is yet another approach and as the name implies it works primarily on the Application Layer of the OSI Model.

Application Layer Firewalls actually, for lack of a better term, understand certain applications and protocols. Examples would be FTP, DNS and HTTP. This form of filtration is able to check to see if the packets are actually behaving incorrectly or if the packets have been incorrectly formatted for the protocol that is indicated. This process also allows for the use of deep packet inspection and the sharing of functionality with Intrusion Prevention Systems (IPS).

Application-layer firewalls work on the application level of the TCP/IP stack (i.e., all browser traffic, or all telnet or ftp traffic), and may intercept all packets traveling to or from an application. They block other packets (usually

dropping them without acknowledgment to the sender). Application firewalls work much like a packet filter but application filters apply filtering rules (allow/block) on a per process basis instead of filtering connections on a per port basis.

On inspecting all packets for improper content, firewalls can restrict or prevent outright the spread of networked computer worms and trojans. The additional inspection criteria can add extra latency to the forwarding of packets to their destination.

Proxy servers

A proxy server is an appliance or application that acts as an intermediary for communicating between computers. A computer has a request for information. The packets are sent to the designated resource but before they can get there they are blocked by the proxy server saying that it will take the request and pass it on. The Proxy Server processes the request and if it is valid it passes onto the designated computer. The designated computer gets the packet and processes the request, sending the answer back to the proxy server. The proxy server sends the information back to the originating computer. It's all a little like a situation with two people who refuse to talk directly with each other using someone else to take messages back and forth.

From a security stand point a Proxy Server can serve a few purposes:

- Protects the anonymity of the originating computer
- The two computers never deal directly with each other
- Packets that are not configured to be forwarded are dropped before reaching the destination computer.
- If malicious code is sent it will affect the Proxy server with out affecting the originating or sending computer.

Proxies can perform a number of roles including:

- Content Filtering
- Caching
- DNS proxy
- Bypassing Filters and Censorship
- Logging and eavesdropping
- Gateways to private networks
- Accessing service anonymously

UTM/ NGFW

Unified Threat Management and Next Generation Firewall are terms originally coined by market research firms and refer to the concept of a comprehensive security solution provided in a single package. It is basically combining of what used to be accomplished by a number of different security technologies all under a single umbrella or in this case, a single device. On the FortiGate firewall this is achieved by the use of Security Profiles and optimized hardware.

In effect it is going from a previous style of firewall that included among its features:

- Gateway Network Firewall
- Routing
- VPN

To a more complete system that includes:

- Gateway Network Firewall
- Routing

- VPN
- Traffic Optimization
- Proxy Services
- Content Filtering
- Application Control
- Intrusion Protection
- Denial of Service Attack Protection
- Anti-virus
- Anti-spam
- Data Leak Prevention
- Endpoint Control of Security Applications
- Load Balancing
- WiFi Access Management
- Authentication Integration into Gateway Security
- Logging
- Reporting

Advantages of using security profiles

- Avoidance of multiple installations.
- Hardware requirements are fewer.
- Fewer hardware maintenance requirements.
- Less space required.
- Compatibility - multiple installations of products increase the probability of incompatibility between systems.
- Easier support and management.
- There is only one product to learn therefore a reduced requirement of technical knowledge.
- Only a single vendor so there are fewer support contracts and Service Level Agreements.
- Easier to incorporate into existing security architecture.
- Plug and play architecture.
- Web based GUI for administration.

FortiGate modes

The FortiGate unit has a choice of modes that it can be used in, either NAT/Route mode or transparent mode. The FortiGate unit is able to operate as a firewall in both modes, but some of its features are limited in transparent mode. It is always best to choose which mode you are going to be using at the beginning of the set up. Once you start configuring the device, if you want to change the mode you are going to lose all configuration settings in the change process.

NAT/Route mode

NAT/Route mode is the most commonly used mode by a significant margin and is thus the default setting on the device. As the name implies the function of NAT is commonly used in this mode and is easily configured but there is no requirement to use NAT. The FortiGate unit performs network address translation before IP packets are sent to the destination network.

These are some of the characteristics of NAT/Route mode:

- Typically used when the FortiGate unit is a gateway between private and public networks.
- Can act as a router between multiple networks within a network infrastructure.
- When used, the FortiGate unit is visible to the networks that it is connected to.
- Each logical interface is on a distinct subnet.
- Each Interface needs to be assigned a valid IP address for the subnet that it is connected to it.

Transparent mode

Transparent mode is so named because the device is effectively transparent in that it does not appear on the network in the way that other network devices show as nodes in the path of network traffic. Transparent mode is typically used to apply the FortiOS features such as Security Profiles etc. on a private network where the FortiGate unit will be behind an existing firewall or router.

These are some of the characteristics of transparent mode:

- The FortiGate unit is invisible to the network.
- All of its interfaces are on the same subnet and share the same IP address.
- The FortiGate unit uses a Management IP address for the purposes of Administration.
- Still able to use NAT to a degree, but the configuration is less straightforward

In transparent mode, you can also perform NAT by creating a security policy or policies that translates the source addresses of packets passing through the FortiGate unit as well as virtual IP addresses and/or IP pools.

How packets are handled by FortiOS

To give you idea of what happens to a packet as it makes its way through the FortiGate unit here is a brief overview. This particular trip of the packet is starting on the Internet side of the FortiGate firewall and ends with the packet exiting to the Internal network. An outbound trip would be similar. At any point in the path if the packet is going through what would be considered a filtering process and if fails the filter check the packet is dropped and does not continue any further down the path.

This information is covered in more detail in other in the Troubleshooting chapter of the FortiOS Handbook in the Life of a Packet section.

The incoming packet arrives at the external interface. This process of entering the device is referred to as **ingress**.

Step #1 - Ingress

1. Denial of Service Sensor
2. IP integrity header checking
3. IPsec connection check
4. Destination NAT
5. Routing

Step #2 - Stateful inspection engine

1. Session Helpers
2. Management Traffic
3. SSL VPN
4. User Authentication

5. Traffic Shaping
6. Session Tracking
7. Policy lookup

Step #3 - Security profiles scanning process

1. Flow-based Inspection Engine
2. IPS
3. Application Control
4. Data Leak Prevention
5. Email Filter
6. Web Filter
7. Anti-virus
8. Proxy-based Inspection Engine
9. VoIP Inspection
10. Data Leak Prevention
11. Email Filter
12. Web Filter
13. Anti-virus
14. ICAP

Step #4 - Egress

1. IPsec
2. Source NAT
3. Routing

Interfaces and zones

A Firewall is a gateway device that may be the nexus point for more than 2 networks. The interface that the traffic is coming in on and should be going out on is a fundamental concern for the purposes of routing as well as security. Routing, policies and addresses are all associated with interfaces. The interface is essentially the connection point of a subnet to the FortiGate unit and once connected can be connected to other subnets.

The following types of interfaces are found on a FortiGate:

- Interface , this can refer to a physical or virtual interface
- Zone
- Virtual Wired Pair

Interfaces

Physical interfaces or not the only ones that need to be considered. There are also virtual interfaces that can be applied to security policies. VLANs are one such virtual interface. Interfaces if certain VPN tunnels are another.

Policies are the foundation of the traffic control in a firewall and the Interfaces and addressing is the foundation that policies are based upon. Using the identity of the interface that the traffic connects to the FortiGate unit tells the firewall the initial direction of the traffic. The direction of the traffic is one of the determining factors in

deciding how the traffic should be dealt with. You can tell that interfaces are a fundamental part of the policies because, by default, this is the criteria that the policies are sorted by.

Zones

Zones are a mechanism that was created to help in the administration of the firewalls. If you have a FortiGate unit with a large number of ports and a large number of nodes in your network the chances are high that there is going to be some duplication of policies. Zones provide the option of logically grouping multiple virtual and physical FortiGate firewall interfaces. The zones can then be used to apply security policies to control the incoming and outgoing traffic on those interfaces. This helps to keep the administration of the firewall simple and maintain consistency.

For example you may have several floors of people and each of the port interfaces could go to a separate floor where it connects to a switch controlling a different subnet. The people may be on different subnets but in terms of security they have the same requirements. If there were 4 floors and 4 interfaces a separate policy would have to be written for each floor to be allowed out on to the Internet off the WAN1 interface. This is not too bad if that is all that is being done, but now start adding the use of more complicated policy scenarios with Security Profiles, then throw in a number of Identity based issues and then add the complication that people in that organization tend to move around in that building between floors with their notebook computers.

Each time a policy is created for each of those floors there is a chance of an inconsistency cropping up. Rather than make up an additional duplicate set of policies for each floor, a zone can be created that combines multiple interfaces. And then a single policy can be created that uses that zone as one side of the traffic connection.



You cannot add physical interfaces associated with VLANs to zones.

Virtual wire pair

The simplified explanation is that two interfaces are set up so that whatever traffic goes through one of the pair is replicated on the other. They are most commonly used when scanning is needed on an interface without interfering with the traffic. On interface "A", everything goes through unaffected. The replicated traffic on interface "B" is sent to an analyzer of some kind and the traffic can be thoroughly scanned without worry of impacting performance.

When two physical interfaces are setup as a Virtual Wire Pair, they will have no IP addressing and are treated similar to a transparent mode VDOM. All packets accepted by one of the interfaces in a virtual wire pair can only exit the FortiGate through the other interface in the virtual wire pair and only if allowed by a virtual wire pair firewall policy. Packets arriving on other interfaces cannot be routed to the interfaces in a virtual wire pair. A FortiGate can have multiple virtual wire pairs.

You cannot add VLANs to virtual wire pairs. However, you can enable wildcard VLANs for a virtual wire pair. This means that all VLAN-tagged traffic can pass through the virtual wire pair if allowed by virtual wire pair firewall policies.

Access control lists

Access control lists (ACLs) in the FortiOS firmware could be considered a granular or more specifically targeted blacklist. These ACLs drop IPv4 or IPv6 packets at the physical network interface before the packets are analyzed by the CPU. On a busy appliance this can really help the performance.

The ACL feature is available on FortiGates with NP6-accelerated interfaces. ACL checking is one of the first things that happens to the packet and checking is done by the NP6 processor. The result is very efficient protection that does not use CPU or memory resources.

Incoming interfaces

The configuration of the Access Control List allow you to specify which in interface the ACL will be applied to. There is a hardware limitation that needs to be taken into account. The ACL is a Layer 2 function and is offloaded to the ISF hardware, therefore no CPU resources are used in the processing of the ACL. It is handled by the inside switch chip which can do hardware acceleration, increasing the performance of the FortiGate. The drawback is that the ACL function is only supported on switch fabric driven interfaces. It also cannot be applied to hardware switch interfaces or their members. Ports such as WAN1 or WAN2 that are found on some models that use network cards that connect to the CPU through a PCIe bus will not support ACL.

Addresses

Because the address portion of an entry is based on a FortiGate address object, it can be any of the address types used by the FortiGate, including address ranges. There is further granularity by specifying both the source and destination addresses. The traffic is blocked not on an either or basis of these addresses but the combination of the two, so that they both have to be correct for the traffic to be denied. Of course, If you want to block all of the traffic from a specific address all you have to do is make the destination address "all".

Because the blocking takes place at the interface based on the information in the packet header and before any processing such as NAT can take place, a slightly different approach may be required. For instance, if you are trying to protect a VIP which has an external address of x.x.x.x and is forwarded to an internal address of y.y.y.y, the destination address that should be used is x.x.x.x, because that is the address that will be in the packet's header when it hits the incoming interface.

Services

Further granulation of the filter by which the traffic will be denied is done by specifying which service the traffic will use.

Firewall policies

The firewall policy is the axis around which most of the other features of the FortiGate firewall revolve. A large portion of the settings in the firewall at some point will end up relating to or being associated with the firewall policies and the traffic that they govern. Any traffic going through a FortiGate unit has to be associated with a policy. These policies are essentially discrete compartmentalized sets of instructions that control the traffic flow going through the firewall. These instructions control where the traffic goes, how it's processed, if it's processed and even whether or not it's allowed to pass through the FortiGate.

When the firewall receives a connection packet, it analyzes the packet's source address, destination address, and service (by port number). It also registers the incoming interface, the outgoing interface it will need to use and the time of day. Using this information the FortiGate firewall attempts to locate a security policy that matches the packet. If it finds a policy that matches the parameters it then looks at the action for that policy. If it is **ACCEPT** the traffic is allowed to proceed to the next step. If the Action is **DENY** or a match cannot be found the traffic is not allowed to proceed.

The 2 basic actions at the initial connection are either **ACCEPT** or **DENY**:

- If the **Action** is **ACCEPT**, the policy action permits communication sessions. There may be other packet processing instructions, such as requiring authentication to use the policy or restrictions on the source and

destination of the traffic.

- If the **Action** is **DENY**, the policy action blocks communication sessions, and you can optionally log the denied traffic. If no security policy matches the traffic, the packets are dropped. A **DENY** security policy is needed when it is required to log the denied traffic, also called “violation traffic”.

There are two other Actions that can be associated with the policy:

- **LEARN** - This is a specialized variation on the **ACCEPT** action. That is set up to allow traffic but to keep traffic logs so that the administrator can go through them to learn what kind of traffic has to be dealt with.
- **IPsec** - This is an **ACCEPT** action that is specifically for IPsec VPNs.

There can also be a number of instructions associated with a FortiGate firewall in addition to the **ACCEPT** or **DENY** actions, some of which are optional. Instructions on how to process the traffic can also include such things as:

- Logging Traffic
- Authentication
- Network Address Translation or Port Address Translation
- Use Virtual IPs or IP Pools
- Caching
- Whether the source of the traffic is based on address, user, device or a combination
- Whether to treat as regular traffic or IPsec traffic
- What certificates to use
- Security profiles to apply
- Proxy Options
- Traffic Shaping

Firewall policy parameters

As mentioned before, for traffic to flow through the FortiGate firewall there must be a policy that matches its parameters:

Incoming interface(s)

This is the interface or interfaces that the traffic is first connection to the FortiGate unit by. The exception being traffic that the FortiGate generates itself. This is not limited to the physical Ethernet ports found on the device. The incoming interface can also be a logical or virtual interface such as a VPN tunnel, a Virtual WAN link or a wireless interface.

Outgoing interface(s)

After the firewall has processed the traffic it needs to leave a port to get to its destination and this will be the interface or interfaces that the traffic leaves by. This interface, like the **Incoming Interface** is not limited to only physical interfaces.

Source address(es)

The addresses that a policy can receive traffic from can be wide open or tightly controlled. For a public web server that the world at large should be able to access, the best choice will be “all”. If the destination is a private web server that only the branch offices of a company should be able to access or a list of internal computers that are

the only ones allowed to access an external resource then a group of preconfigured addresses is the better strategy.

Additional parameters under the Source Address, though they are not mandatory are:

- **Source User(s)**

This parameter is based on a user identity that can be from a number of authentication authorities. It will be an account or group that has been set up in advance that can be selected from the drop down menu. The exception to this is the feature that allows the importing of LDAP Users. When the feature is used, a small wizard window will appear to guide the user through the setup. The caveat is that the LDAP server object in the **User and Device > Authentication > LDAP Servers** section has to be already configured to allow the use of this import feature.

- **Source Device Type**

This parameter is for narrowing down the traffic sending devices to those that the FortiGate is familiar with. Again the contents of this parameter need to be a preconfigured object and these are defined at **User and Device > Custom Devices & Groups**. This parameter can limit the devices that can connect to this policy to those specific MAC addresses that are already known by the FortiGate and are approved for the policy.

Destination address(es)

In the same way that the source address may need to be limited, the destination address can be used as a traffic filter. When the traffic is destined for internal resources the specific address of the resource can be defined to better protect the other resources on the network. One of the specialized destination address options is to use a Virtual IP address. The destination address doesn't need to be internal you can define policies that are only for connecting to specific addresses on the Internet.

Internet service(s)

In this context, and Internet service is a combination of one or more addresses and one or more services associated with a service found on the Internet such as an update service for software.

Schedule

The time frame that is applied to the policy. This can be something as simple as a time range that the sessions are allowed to start such as between 8:00 am and 5:00 pm. Something more complex like business hours that include a break for lunch and time of the session's initiation may need a schedule group because it will require multiple time ranges to make up the schedule.

Service

The service or service chosen here represent the TCP/IP suite port numbers that will most commonly be used to transport the named protocols or group of protocols. This will be a little different than Application Control which looks more closely at the packets to determine the actual protocol used to create them.

Without all six (possibly 8) of these things matching, the traffic will be declined. Each traffic flow requires a policy and the direction is important as well. Just because packets can go from point A to point B on port X does not mean that the traffic can flow from point B to point A on port X. A policy must be configured for each direction.

When designing a policy there is often reference to the traffic flow, but most communication is a two way connection so trying to determine the direction of the flow can be somewhat confusing. If traffic is HTTP web traffic the user sends a request to the web site, but most of the traffic flow will be coming from the web site to the user. Is the traffic flow considered to be from the user to the web site, the web site to the user or in both directions? For the purposes of determining the direction for a policy the important factor is the direction of the

initiating communication. The user is sending a request to the web site so this is the initial communication and the web site is just responding to it so the traffic will be from the users network to the Internet.

A case where either side can initiate the communication like between two internal interfaces on the FortiGate unit would be a more likely situation to require a policy for each direction.

Application groups for NGFW policies

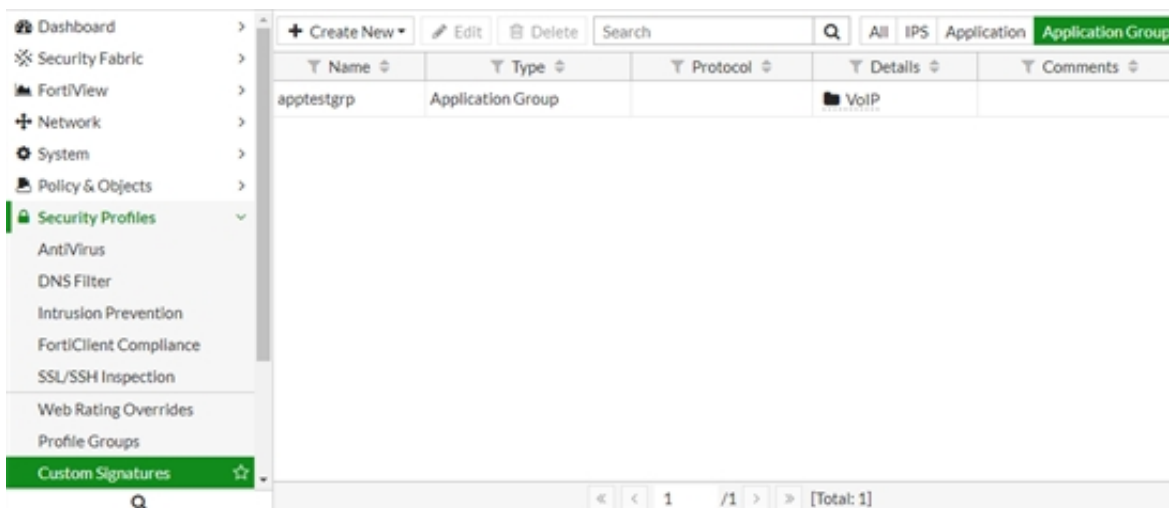
In addition to parameters like schedule and service, NGFW policies can filter by application or application category.



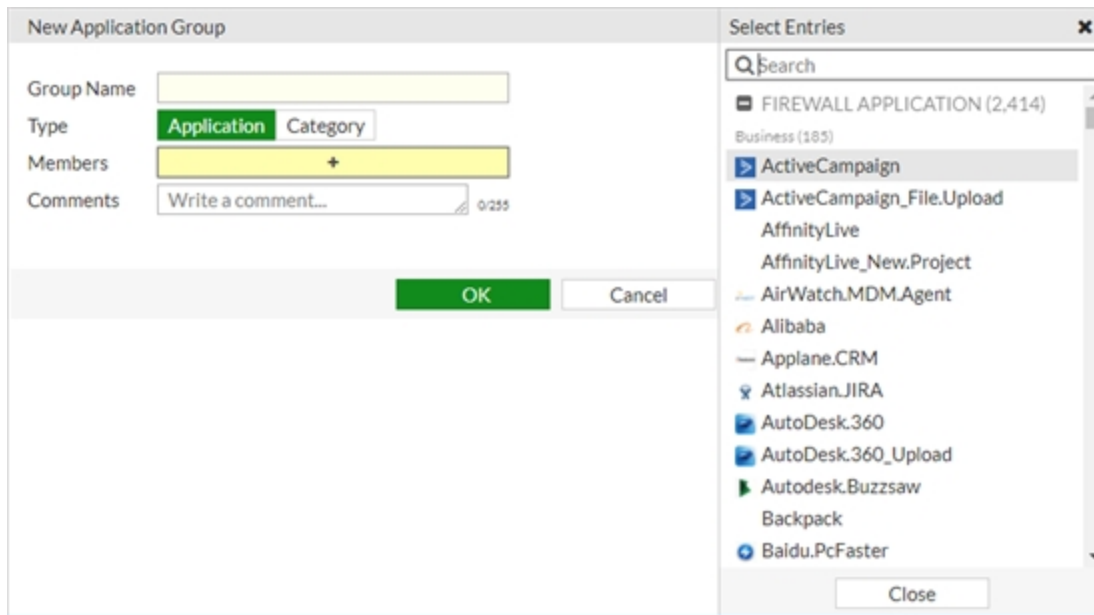
The use of this feature is dependent upon the VDOM having the following **System Operation Settings**:

- **Inspection Mode: Flow-based**
- **NGFW Mode: Policy-based**

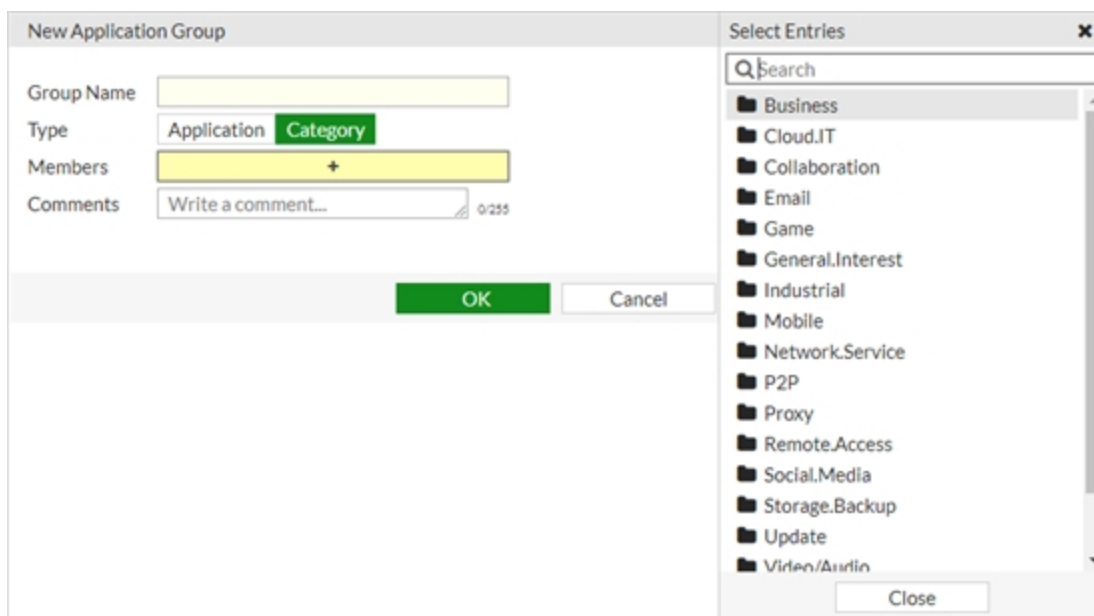
To use the feature first create an application group in **Security Profiles > Custom Signatures**.



From the editing page for the **New Application Group**, choose a group type of **Application** and select individual applications for membership in the group.



Alternatively, select **Category** and add one or more application categories as group members.



Whichever type of **Application Group** you choose, the available **Members** will be displayed in the selection pane that slides out from the right of the window.

Once the **Application Group** is created, you can apply it to a policy in the **Application** field, by clicking on the **+** in the field and selecting members from the options under the **Group** tab at the top of the pane that slides out from the right of the window.



CLI

To create or edit an application group:

```
config application group
  edit <group_name>
    set comments
    set type {application | category}
    set application <Application ID number>
    set category <category ID number>
  end
```

To add an application group to a policy:

```
config firewall policy
  edit 1
    set app-group "test" "test1"
  end
```

Application ID number

In the CLI, you add applications to a group by using the application ID number. To see the list of application ID numbers, run the following command when `type` is set to `application`:

```
set application ? <enter>
```

The start of the list looks like:

```
set application
ID          Select Application ID
38614       1kxun
29025       1undl.Mail
17534       2ch
17535       2ch_Post
16284       3PC
16616       4shared
```

```

35760      4shared_File.Download
34742      4shared_File.Upload
44606      5ch
44607      5ch_Post
38923      8tracks
17045      9PFS
16554      126.Mail
23345      360.Safeguard.Update
35963      360.Yunpan
35967      360.Yunpan_File.Download
35966      360.Yunpan_File.Upload
42324      360.Yunpan_Login
16413      A.N
31529      ABC
...

```

Only the first 20 have been listed here.

Category ID number

The ID numbers for the categories in the CLI are found in the same manner as the applications. When the `type` is set to `category`, run the command:

```
set category ? <enter>
```

This list is shorter.

```

set category
ID          Select Category ID
2           P2P
3           VoIP
5           Video/Audio
6           Proxy
7           Remote.Access
8           Game
12          General.Interest
15          Network.Service
17          Update
21          Email
22          Storage.Backup
23          Social.Media

```

25	Web.Client
26	Industrial
28	Collaboration
29	Business
30	Cloud.IT
31	Mobile

What is not expressly allowed is denied

One of the fundamental ideas that can be found in just about any firewall is the rule than anything that is not expressly allowed is by default denied. This is the foundation for any strategy of protecting your network. Right out of the box, once you have your FortiGate device connected into your network and hooked up with your ISP, your network is protected. Nothing is getting out or in so it is not very convenient, but you don't have to worry that between the time you hooked it up and the point that you got all of the policies in place that someone could have gotten in and done something to your resources. The reason that this needs to be kept in mind when designing policies is because you cannot assume that any traffic will be allowed just because it makes sense to do so. If you want any kind of traffic to make it past the FortiGate firewall you need to create a policy that will allow that traffic. To maintain the protection of the network should also make sure that the any policy you create allows only the traffic you intend to go only to where you specifically want it to go and when you want it to go there.

Example

You have a web server on your network that is meant to provide a collaborative work environment web site for your employees and a partner company for a project over the course of the next 3 months.

It is theoretically possible to allow connections into your network to any device on that network for any service and at any time. The problem with this is that we might not want just anybody looking at those resources. Sadly, no matter how much it is wished otherwise, not everybody on the Internet can be trusted. Which means we now have to be very specific in our instructions as to what traffic to allow into the network. Each step that we take towards being more specific as to what we allow means that there is that much more that is not allowed and the level of protection of a resources is directly proportional to the amount of traffic that is not allowed. If somebody can't get at it they can't damage or steal it.

Limiting where the traffic is allowed to go to means that other computers on your network besides the web-server are protected.

- Limiting where the traffic is allowed to come from means that, if feasible, you can limit the systems that can access the web server to just employees or the partner company computers.
- Limiting the services to just web traffic means that a malicious person, even if they were connection from a computer at the partner organization could only use the features of web traffic to do anything malicious.
- Limiting the policy to the time span of the project would mean that even if the IT department forgot to remove the policy after the end of the project than no computer from the other company could be used to do anything malicious through the policy that allowed the traffic.

This is just a very basic example but it shows the underlying principles of how the idea that anything not expressly allowed is by default denied can be used to effectively protect your network.

Policy order

Another important factor in how firewall policies work is the concept of precedence of order or if you prefer a more recognizable term, “first come, first served”.

It is highly likely that even after only a relatively small number of policies have been created that there will be some that overlap or are subsets of the parameters that the policies use to determine which policy should be matched against the incoming traffic. When this happens there has to be a method to determine which policy should be applied to the packet. The method which is used by most firewalls is based on the order of the sequence of the policies.

If all of the policies were placed in a sequential list the process to match up the packet would start at the top of the list and work its way down. It would compare information about the packet, specifically these points of information:

1. The interface the packet connected to the FortiGate firewall
2. The source of the packet. This can include variations of the address, user credentials or device
3. The destination of the packet. This can include address or Internet service
4. The interface the packet would need to use to get to the destination address based on the routing table
5. The service or port the packet is destined for
6. The time that the packet connected to the FortiGate

As soon as the a policy is reached that matches all of the applicable parameters, the instructions of that policy are applied and the search for any other matching policies is stopped. All subsequent policies are disregarded. Only 1 policy is applied to the packet.

If there is no matching policy among the policies that have been configured for traffic the packet finally drops down to what is always the last policy. It is an implicit policy. One of a few that are referred to by the term “policy0”. This policy denies everything.

The implicit policy is made up of the following settings:

- Incoming Interface: any
- Source Address: any
- Outgoing Interface: any
- Destination Address: any
- Action: DENY

The only setting that is editable in the implicit policy is the logging of violation traffic.

A logical best practice that comes from the knowledge of how this process works is to make sure that the more specific or specialized a policy is, the closer to the beginning of the sequence it should be. The more general a policy is the higher the likelihood that it could include in its range of parameters a more specifically targeted policy. The more specific a policy is, the higher the probability that there is a requirement for treating that traffic in a specific way.

Example

For security reasons there is no FTP traffic allowed out of a specific subnet so there is a policy that states that any traffic coming from that subnet is denied if the service is FTP, so the following policy was created:

Policy #1

Source Interface	Internal1
Source Address	192.168.1.0/24
Source User(s)	<left at default setting>
Source Device Type	<left at default setting>
Outgoing Interface	WAN1
Destination Address	0.0.0.0/0.0.0.0
Service	FTP
Schedule	always
Action	deny

Now as these things usually go it turns out that there has to be an exception to the rule. There is one very secure computer on the subnet that is allowed to use FTP and once the content has been checked it can then be distributed to the other computer on the subnet. So a second firewall policy is created.

Policy #2

Source Interface	Internal1
Source Address	192.168.1.38/32
Source User(s)	<left at default setting>
Source Device Type	<left at default setting>
Outgoing Interface	WAN1
Destination Address	0.0.0.0/0.0.0.0
Service	FTP
Schedule	always
Action	Allow

By default, a policy that has just been created will be placed last in the sequence so that it is less likely to interfere with existing policies before it can be moved to its intended position. If you look at Policy #2 you will notice that it is essentially the same as Policy #1 exempt for the Source Address and the Action. You will also notice that the Source Address of the Policy #2 is a subset of the Source address in policy #1. This means that if nothing further is done, Policy #2 will never see any traffic because the traffic will always be matched by Policy #1 and processed before it has a chance to reach the second policy in the sequence. For both policies to work as intended Policy #2 needs to be moved to before Policy #1 in the sequence.

Policy identification

There are two ways to identify a policy. The most obvious is the policy name and this is easily read by humans, but with a little effort it is possible to have a policy without a name, therefore every policy has an ID number.

When looking at the policy listing it can appear as if the policies are identified by the sequence number in the far left column. The problem is that this number changes as the position of the policy in the sequence changes. The column that correctly identifies the policy, and the value sticks with the policy is the "ID" column. This column is not shown by default in the listing but can be added to the displayed columns by right clicking on the column heading bar and selecting it from the list of possible columns.

When looking in the configuration file the sequence is based upon the order of the policies as they are in the file just as they are in the list in the GUI. However, if you need to edit the policy in the CLI you must use the ID number.

UUID support

Universally Unique Identifier (UUID) attributes have been added to policies to improve functionality when working with FortiManager or FortiAnalyzer units. If required, the UUID can be set manually through the CLI.

CLI Syntax:

```
config firewall {policy/policy6/policy46/policy64}
  edit 1
    set uuid <example uuid: 8289ef80-f879-51e2-20dd-fa62c5c51f44>
  next
end
```

NTurbo support CAPWAP traffic

NTurbo is used for IPSec+IPS case. The IPSec SA info is passed to NTurbo as part of VTAG for control packet and will be used for the xmit.



If the packets need to go through IPSec interface, the traffic will be always offloaded to NTurbo. But for the case that SA has not been installed to NP6 because of hardware limitation or SA offload disable, the packets will be sent out through raw socket by IPS instead of NTurbo, since the software encryption is needed in this case.

CLI :

Previously, NTurbo could only be enabled or disabled globally. The setting of np-acceleration has been added to the firewall policy context instead of just the global context.

CLI command in the firewall policy to enable/disable NTurbo acceleration.

```
config firewall policy
  edit 1
    set np-acceleration [enable|disable]
```

end

When IPS is enabled for VPN IPsec traffic, the data can be accelerated by NTurbo.

Learning mode for policies

The learning mode feature is a quick and easy method for setting a policy to allow everything but to log it all so that it can later be used to determine what restrictions and protections should be applied. The objective is to monitor the traffic not act upon it while in Learning mode.

Once the **Learn** action is enabled, functions produce hard coded profiles that will be enabled on the policy. The following profiles are set up:

- AntiVirus (av-profile)
- Web Filter (webfilter-profile)
- Anti Spam(spamfilter-profile)
- Data Leak Prevention (dlp-sensor)
- Intrusion Protection (ips-sensor)
- Application Control (application-list)
- Proxy Options (profile-protocol-options)



- These UTM profiles are all using Flow mode
 - SSL inspection is always disabled for the Learn option
 - These profiles are static and cannot be edited.
-

Profiles that are not being used are:

- DNS Filter (Does not have a Flow mode)
- Web Application Firewall (Does not have a Flow mode)
- CASI (Almost all signatures in CASI require SSL deep inspection. Without SSL inspection, turning on CASI serves little purpose)

The ability to allow policies to be set to a learning mode is enabled on a per VDOM basis.

```
config system settings
  set gui-policy-learning [enable | disable]
end
```

Once the feature is enabled on the VDOM, Learn is an available **Action** option when editing a policy.



Because this feature requires a minimum level of logging capabilities, it is only available on FortiGates with hard drives. Smaller models may not be able to use this feature.

New Policy

Name

Incoming Interface

+

Outgoing Interface

+

Source

+

Destination Address

+

Schedule

always

Service

+

Action

✓ ACCEPT

✗ DENY

LEARN

IPsec

Firewall / Network Options

NAT

☒

Comments

Write a comment...

0/1023

Enable this policy

☒

OK

Cancel

Once the Learning policy has been running for a sufficient time to collect needed information a report can be looked at by going to **Log & Report > Learning Report**.

The Report can be either a **Full Report** or a **Report Summary**

The time frame of the report can be **5 minutes**, **1 hour**, or **24 hours**.

The Learning Report includes:

Deployment Methodology

- Test Details
 - Start time
 - End time
 - Model
 - Firmware
- Policy List

Executive Summary

- Total Attacks Detected
- Top Application Category
- Top Web Category
- Top Web Domain
- Top Host by Bandwidth
- Host with Highest Session Count

Security and Threat Prevention

- High Risk Applications
- Application Vulnerability Exploits

- Malware, botnets and Spyware/Adware
- At-Risk Devices and Hosts

User Productivity


- Application Usage
 - Top Application Categories
 - Top Social Media Applications
 - Top Video/Audio Streaming Applications
 - Top Peer to Peer Applications
 - Top Gaming Applications
- Web Usage
 - Top Web Categories
 - Top Web Applications
 - Top Web Domains

Policy modes

You can operate your FortiGate or individual VDOMs in **Next Generation Firewall (NGFW) Policy Mode**.

You can enable NGFW policy mode by going to **System > Settings**, setting the **Inspection mode** to **Flow-based** and setting the NGFW mode to **Policy-based**. When selecting **NGFW policy-based** mode you also select the SSL/SSH Inspection mode that is applied to all policies

Flow-based inspection with profile-based **NGFW mode** is the default in FortiOS 5.6.







Or use the following CLI command:

```
config system settings
  set inspection-mode flow
  set ngfw-mode {profile-based | policy-based}
end
```

NGFW policy mode and NAT

If your FortiGate is operating in NAT mode, rather than enabling source NAT in individual NGFW policies you go to **Policy & Objects > Central SNAT** and add source NAT policies that apply to all matching traffic. In many cases you may only need one SNAT policy for each interface pair. For example, if you allow users on the internal network (connected to port1) to browse the Internet (connected to port2) you can add a port1 to port2 Central SNAT policy similar to the following:

New Central SNAT Policy




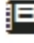











Incoming Interface	<div> port1</div> <div>+</div> <div>✕</div>
Outgoing Interface	<div> port2</div> <div>+</div> <div>✕</div>
Source address	<div> all</div> <div>+</div> <div>✕</div>
Destination address	<div> all</div> <div>+</div> <div>✕</div>

☒ NAT

IP Pool Configuration	<div>Use Outgoing Interface Address</div> <div>Use Dynamic IP Pool</div>
Protocol	<div>ANY</div> <div>TCP</div> <div>UDP</div> <div>SCTP</div> <div>Specify</div> <div>0</div>











Application control in NGFW policy mode

You configure **Application Control** simply by adding individual applications to security policies. You can set the action to accept or deny to allow or block the applications.

Name 	Block YouTube	
Incoming Interface	 port1	▼
Outgoing Interface	 port2	▼
Source	 all	✕
	+	
Destination	 all	✕
	+	
Schedule	 always	▼
Service	 ALL	✕
	+	
Application	 YouTube ✕  YouTube_Channel.Access ✕  YouTube_HD.Streaming ✕  YouTube_Video.Access ✕  YouTube_Video.Embedded ✕ +	
URL Category	+	
Action	 ACCEPT  DENY  LEARN	

Web filtering in NGFW mode

You configure **Web Filtering** by adding URL categories to security policies. You can set the action to accept or deny to allow or block the applications.

Name 	Block Streaming Websites	
Incoming Interface	 port1	▼
Outgoing Interface	 port2	▼
Source	 all	✕
	+	
Destination	 all	✕
	+	
Schedule	 always	▼
Service	 ALL	✕
	+	
Application	+	
URL Category	Streaming Media and Download	✕
	+	
Action	 ACCEPT  DENY  LEARN	

Other NGFW policy mode options

You can also combine both application control and web filtering in the same NGFW policy mode policy. Also if the policy accepts applications or URL categories you can also apply Antivirus, DNS Filtering, and IPS profiles in NGFW mode policies as well a logging and policy learning mode.

Security profiles

Where security policies provide the instructions to the FortiGate unit for controlling what traffic is allowed through the device, the Security profiles provide the screening that filters the content coming and going on the network. Security profiles enable you to instruct the FortiGate unit about what to look for in the traffic that you don't want, or want to monitor, as it passes through the device.

A security profile is a group of options and filters that you can apply to one or more firewall policies. Security profiles can be used by more than one security policy. You can configure sets of security profiles for the traffic types handled by a set of security policies that require identical protection levels and types, rather than repeatedly configuring those same security profile settings for each individual security policy.

For example, while traffic between trusted and untrusted networks might need strict antivirus protection, traffic between trusted internal addresses might need moderate antivirus protection. To provide the different levels of protection, you might configure two separate profiles: one for traffic between trusted networks, and one for traffic between trusted and untrusted networks.

Security profiles are available for various unwanted traffic and network threats. Each are configured separately and can be used in different groupings as needed. You configure security profiles in the Security Profiles menu and applied when creating a security policy by selecting the security profile type.

There is a separate handbook for the topic of the Security Profiles, but because the Security Profiles are applied through the Firewall policies it makes sense to have at least a basic idea of what the security profile do and how they integrate into the FortiGate's firewall policies. The following is a listing and a brief description of what the security profiles offer by way of functionality and how they can be configured into the firewall policies.

- HTTP
- SMTP
- POP3
- IMAP
- FTP
- NNTP
- MAPI
- DNS
- IM

AntiVirus

Antivirus is used as a catch all term to describe the technology for protection against the transmission of malicious computer code sometimes referred to as malware. As anyone who has listened to the media has heard that the Internet can be a dangerous place filled with malware of various flavors. Currently, the malware that is most common in the Internet, in descending order, is Trojan horses, viruses, worms, adware, back door exploits, spyware and other variations. In recent years, not only has the volume of malicious software become greater than would have been believed when it first appeared but the level of sophistication has risen as well.

The Antivirus Filter works by inspecting the traffic that is about to be transmitted through the FortiGate. To increase the efficiency of effort it only inspects the traffic being transmitted via the protocols that it has been configured to check. Before the data moves across the FortiGate firewall from one interface to another it is checked for attributes or signatures that have been known to be associated with malware. If malware is detected, it is removed.

Web Filtering

Malicious code is not the only thing to be wary of on the Internet. There is also the actual content. While the content will not damage or steal information from your computer there is still a number of reasons that would require protection from it.

In a setting where there are children or other sensitive people using the access provided by a connected computer there is a need to make sure that images or information that is not appropriate is not inadvertently displayed to them. Even if there is supervision, in the time it takes to recognize something that is inappropriate and then properly react can expose those we wish to protect. It is more efficient to make sure that the content cannot reach the screen in the first place.

In an organizational setting, there is still the expectation that organization will do what it can to prevent inappropriate content from getting onto the computer screens and thus provoking an Human Resources incident. There is also the potential loss of productivity that can take place if people have unfiltered access to the Internet. Some organizations prefer to limit the amount of distractions available to tempt their workers away from their duties.

The Web filter works primarily by looking at the destination location request for a HTTP(S) request made by the sending computer. If the URL is on a list that you have configured to list unwanted sites, the connection will be disallowed. If the site is part of a category of sites that you have configured to deny connections to the session will also be denied. You can also configure the content filter to check for specific key strings of data on the actual web site and if any of those strings of data appear the connection will not be allowed.

The configuration for each of these protocols is handled separately.

DNS filtering is similar to Web Filtering from the viewpoint of the user. The difference is under the hood. When using regular Web Filtering, the traffic can go through some processing steps before it gets to the point where the web filter determines whether or not the traffic should be accepted or denied. Because the filtering takes place at the DNS level, some sites can be denied before a lot of the additional processing takes place. This can save resource usage on the FortiGate and help performance.

Application Control

Application Control is designed to allow you to determine what applications are operating on your network and to also filter the use of these applications as required. Application control is also for outgoing traffic to prevent the use of applications that are against an organization's policy from crossing the network gateway to other networks. An example of this would be the use of proxy servers to circumvent the restrictions put in place using the Web Filtering.

Intrusion Protection (IPS)

Intrusion Prevention System is almost self explanatory. In the same way that there is malware out on the Internet that the network needs to be protected from there are also people out there that take a more targeted approach to malicious cyber activity. No operating system is perfect and new vulnerabilities are being discovered all of the time. An intrusion prevention system is designed to look for activity or behavior that is consistent with attacks against your network. When attack like behavior is detected it can either be dropped or just monitored depending on the approach that you would like to take.

As new vulnerabilities are discovered they can be added to the IPS database so that the protection is current.

Anti-Spam

Spam or unsolicited bulk email is said to account for approximately 90% of the email traffic on the Internet. Sorting through it is both time consuming and frustrating. By putting an email filter on policies that handle email traffic, the amount of spam that users have to deal with can be greatly reduced.

Data Leak Prevention (DLP)

Data Leak Prevention is used to prevent sensitive information from leaving your network. When people think of security in the cyber-world one of the most common images is that of a hacker penetrating your network and making off with your sensitive information, but the other way that you can lose sensitive data is if someone already on the inside of your network sends it out. This does not have to be an act of industrial espionage. It can just be a case of not knowing the policies of the organization or a lack of knowledge of security or laws concerning privacy.

For instance, a company may have a policy that they will not reveal anyone's Social Security number, but an employee emails a number of documents to another company that included a lengthy document that has a Social Security number buried deep within it. There is not malicious intent but if the information got out there could be repercussions.

If an organization has any information in a digital format that it cannot afford for financial or legal reasons, to leave its network, it makes sense to have Data Leak Prevention in place as an additional layer of protection.

VoIP

Voice over IP is essentially the protocols for transmitting voice or other multimedia communications over Internet Protocol networks such as the Internet. The Security Profiles VoIP options apply the SIP Application Level Gateway (ALG) to support SIP through the FortiGate unit. The SIP ALG can also be used to protect networks from SIP-based attacks.

ICAP

Internet Content Adaptation Protocol (ICAP) off loads HTTP traffic to another location for specialized processing. The purpose of this module when triggered is to send the incoming HTTP traffic over to a remote server to be processed thus taking some of the strain off of the resources of the FortiGate unit. The reasons for the specialized process could be anything from more sophisticated Antivirus to manipulation of the HTTP headers and URLs.

Just like other components of the FortiGate, there is the option for different Proxy Option profiles so that you can be very granular in your control of the workings of the FortiGate. In the case of the Proxy Option profiles the thing that you will want to focus on is the matching up of the correct profile to a firewall policy that is using the appropriate protocols. If you are creating a Proxy Option profile that is designed for policies that control SMTP traffic into your network you only want to configure the settings that apply to SMTP. You do not need or want to configure the HTTP components.

The Web Application Firewall performs a similar role as devices such as Fortinet's FortiWeb, though in a more limited fashion. Its function is to protect internal web servers from malicious activity specific to those types of servers. This includes things like SQL injection, Cross site Scripting and trojans. It uses signatures and other straight forward methods to protect the web servers, but it is a case of turning the feature on or off and the actions are limited to **Allow**, **Monitor** or **Block**. To get protection that is more sophisticated, granular and intelligent, as well as having many more features, it is necessary to get a device like the FortiWeb that can devote more resources to the process. However, if your needs are simple, choosing to use the WAF feature built into the FortiGate should provide valuable protection.

The comfort client feature mitigates this potential issue by feeding a trickle of data while waiting for the scan to complete so as to let the user know that processing is taking place and that there hasn't been a failure in the transmission. This slow transfer rate continues until the antivirus scan is complete. Once the file has been successfully scanned without any indication of viruses the transfer will proceed at full speed.

Security profile groups

It may seem counter intuitive to have a topic on security profile groups in the Firewall Chapter/Handbook when there is already a chapter/handbook on Security Profiles, but there are reasons.

- Security profile groups are used exclusively in the configuration of a firewall policy, which is described in the Firewall Chapter/Handbook.
- The CLI commands for creating and using security profile groups are in the firewall configuration context of the command line structure of settings.

The purpose of security profile groups is just the same as other groups such as Address, Service, and VIP groups. They are used to save time and effort in the administration of the FortiGate when there are a lot of policies with a similar pattern of Security Profile use. In a fairly basic network setup with a handful of policies it doesn't seem like it would be worth the effort to set up groups of security profiles but if you have a large complex configuration with

hundreds of policies where many of them use the same security profiles it can definitely save some effort and help prevent missing adding an important profile from a policy. As an added benefit, when it comes time to add or change the profiles for the policies that use the Security Profile Groups, the changes only have to be made to the group, not each policy.

The most difficult part about using security profile groups is making them visible in the GUI.

Making security profile groups visible in the GUI

By default, the Security Profile Groups are not visible in the GUI. Neither the ability to assign one to a policy nor the ability to configure the members of a group are available by default. You will not find the option to enable Security Profile Groups under **System > Feature Visibility** either. Instead, they only become visible in the GUI once one has been created and assigned to a policy. This must be done the first time through the CLI using the following syntax:

```
config system settings
    set gui-dynamic-profile-display enable
end
```

Step 1 - Create a security profile group:

Enter the command:

```
config firewall profile-group
```

Use the edit command to give a name to and create a new Security Profile Group

```
(profile-group) # edit test-group
```

Configure the members of the group by setting the name of the desired profile in the field for the related profile/sensor/list. The options are:

av-profile	Name of an existing Antivirus profile.
webfilter-profile	Name of an existing Web filter profile.
dnsfilter-profile	Name of an existing DNS filter profile.
spamfilter-profile	Name of an existing Spam filter profile.
dlp-sensor	Name of an existing DLP sensor.
ips-sensor	Name of an existing IPS sensor.
application-list	Name of an existing Application list.
voip-profile	Name of an existing VoIP profile.
icap-profile	Name of an existing ICAP profile.
waf-profile	Name of an existing Web application firewall profile.
profile-protocol-options	Name of an existing Protocol options profile.
ssl-ssh-profile	Name of an existing SSL SSH profile.

Example:

```
set av-profile default
set profile-protocol-options default
end
```



Always set the `profile-protocol-options` setting before attempting to save the profile group. If this is not set, you will get the error:

```
node_check_object fail! for profile-protocol-options
Attribute 'profile-protocol-options' MUST be set.
Command fail. Return code -56
```

Step 2 - Add a security profile to a policy

Now that there is group to add to a policy we can configure a policy to allow the use of a Security Policy group. This is also done in the CLI.

In the following example only the command necessary to enable the use and pick of a Security Policy group have been listed.

```
config firewall policy
edit 0
set utm-status enable
set profile-type group
set profile-group test-group
```

Step 3 - The appearance in the GUI of the security profile group configuration features

- Under **Security Profiles** there is a menu item called **Profile Groups** that can be used to create new and edit existing profile groups.
- In the **Edit Policy** window for **IPv4** and **IPv6** policies there is a **Use Security Profile Group** field to enable or disable the use of the groups.
 - In the window, policy groups can be created or edited by clicking on the appropriate icons next to or in the drop down menu
- In the policy listing window there is a Security Profiles column.
 - Right or left clicking on the icon for the group brings up editing options either via a slide out window or a drop down menu, respectively.

Proxy option components

Any time a security profile that requires the use of a proxy is enabled the Proxy Options field will be displayed. Certain inspections defined in security profiles require that the traffic be held in proxy while the inspection is carried out and so the Proxy Options are there to define the parameters of how the traffic will be processed and to what level the traffic will be processed. In the same way that there can be multiple security profiles of a single type there can also be a number of unique Proxy Option profiles so that as the requirements for a policy differ from one policy to the next you can also configure a different Proxy Option profile for each individual policy or you can use one profile repeatedly.

The Proxy Options refer to the handling of the following protocols:

- HTTP
- SMTP
- POP3
- IMAP
- FTP
- NNTP
- MAPI
- DNS
- IM

The configuration for each of these protocols is handled separately.

It should also be noted that these configurations apply to only the Security Profiles Proxy-based processes and not the Flow-based processes.

The use of different proxy profiles and profile options

Just like other components of the FortiGate, there is the option for different Proxy Option profiles so that you can be very granular in your control of the workings of the FortiGate. In the case of the Proxy Option profiles the thing that you will want to focus on is the matching up of the correct profile to a firewall policy that is using the appropriate protocols. If you are creating a Proxy Option profile that is designed for policies that control SMTP traffic into your network you only want to configure the settings that apply to SMTP. You do not need or want to configure the HTTP components.

Oversized file log

This setting is for those that would like to log the occurrence of oversized files being processed. It does not change how they are processed it only enables the FortiGate unit to log that they were either blocked or allowed through. A common practice is to allow larger files through without antivirus processing. This allows you to get an idea of how often this happens and decide on whether or not to alter the settings relating to the treatment of oversized files.

The setting of the threshold for what is considered to be an oversized file is located in the Oversized File / Email Threshold that is found in some of the protocol options for the Proxy Options.

Protocol port mapping

While each of the protocols listed has a default TCP port that is commonly used, the level of granularity of control on the FortiGate firewall allows that the port used by the protocols can be individually modified in each separate Profile. It can also be set to inspect any port with flowing traffic for that particular protocol. The headers of the packets will indicate which protocol generated the packet. To optimize the resources of the unit the mapping and inspection of protocols can be enabled or disabled depending on your requirements.

Comfort clients

When proxy-based antivirus scanning is enabled, the FortiGate unit buffers files as they are downloaded. Once the entire file is captured, the FortiGate unit begins scanning the file. During the buffering and scanning procedure, the user must wait. After the scan is completed, if no infection is found, the file is sent to the next step in the process flow. If the file is a large one this part of the process can take some time. In some cases enough time that some users may get impatient and cancel the download.

The comfort client feature mitigates this potential issue by feeding a trickle of data while waiting for the scan to complete so as to let the user know that processing is taking place and that there hasn't been a failure in the transmission. This slow transfer rate continues until the antivirus scan is complete. Once the file has been successfully scanned without any indication of viruses the transfer will proceed at full speed.

If there is evidence of an infection the FortiGate unit caches the URL and drops the connection. The client does not receive any notification of what happened because the download to the client had already started. Instead, the download stops and the user is left with a partially downloaded file. If the user tries to download the same file again within a short period of time, the cached URL is matched and the download is blocked. The client receives the Infection cache message replacement message as a notification that the download has been blocked. The number of URLs in the cache is limited by the size of the cache.

Client comforting is available for HTTP and FTP traffic. If your FortiGate unit supports SSL content scanning and inspection, you can also configure client comforting for HTTPS and FTPS traffic.



Buffering the entire file allows the FortiGate unit to eliminate the danger of missing an infection due to fragmentation because the file is reassembled before examination. Client comforting can send unscanned and therefore potentially infected content to the client. You should only enable client comforting if you are prepared to accept this risk. Keeping the client comforting interval high and the amount low will reduce the amount of potentially infected data that is downloaded.

Oversized file/email threshold

This is another feature that is related to antivirus scanning. The FortiGate unit has a finite amount of resources that can be used to buffer and scan a file. If a large file such as an ISO image or video file was to be downloaded this could not only overwhelm the memory of the FortiGate, especially if there were other large files being downloaded at the same time, but could exceed it as well. For this reason, how to treat large files needs to be addressed.

A threshold is assigned to determine what should be considered an oversized file or email. This can be set at any size from 1 MB to 50 MB. Any file or email over this threshold will not be processed by the Antivirus Security Profiles. Once a file is determined to be oversized it must be then determined whether to allow it or to block it.

These settings are not a technical decision but a policy one that will depend on your comfort level with letting files into your network. As there often is, there is a compromise between convenience or ease of use and security. If you want to go for a high peace of mind level you can configure the firewall to block oversized files and thus no files would be coming into the network that have not been scanned. If you are looking for optimizing the memory of the FortiGate unit and making sure that everybody is getting the files they want, you can lower the threshold and allow files that are over the threshold.



It should be noted that in terms of probability that malware is more likely to be found in smaller files than in larger files. A number of administrators take this into account when they lower the default threshold so as to lessen the impact on memory if they see the FortiGate unit going into conserve mode on a regular basis.

Chunked bypass

The HTTP section allows the enabling of "Chunked Bypass". This refers to the mechanism in version 1.1 of HTTP that allows a web server to start sending chunks of dynamically generated output in response to a request before actually knowing the actual size of the content. Where dynamically generated content is concerned this means

that there is a faster initial response to HTTP requests. From a security stand point it means that the content will not be held in the proxy as an entire file before proceeding.

Allow fragmented messages

The specifications of RFC 2046 allow for the breaking up of emails and sending the fragments in parallel to be rebuilt and read at the other end by the mail server. It was originally designed to increase the performance over slower connections where larger email messages were involved. It will depend on your mail configuration if this is even possible for your network but outside of Microsoft Outlook and Outlook Express, not many email clients are set up to break up messages like this. The drawback of allowing this feature is that if malware is broken up between multiple fragments of the message the risk is run that it will not be detected by some antivirus configurations because the code may not all be present at the same time to identify.

Append email signature

The Append Email Signature is used when an organization would like to ensure that over and above our in this case underneath the existing personal signatures of the sender, all of the emails going out of their network have the appropriate “boilerplate”, for lack of a better term. These appended emails do not replace existing signatures. They are as the feature states, appended to the email.

Examples could include things like:

- Without prior approval the email should not be forwarded.
- Please be environmentally friendly and don't print out emails
- For questions regarding the purchasing of our products please call...

It can be anything that the organization would like as long as it is in text format. The use of this feature usually works best in an environment where there is some standardization of what goes into the personal signatures of the senders so that there is no duplication or contradiction of information in the signatures.

SSL/SSH inspection

While the profile configuration for **SSL/SSH Inspection** is found in the **Security Profiles** section it is enabled in the firewall policy by enabling any of the security profiles. Choosing which of the **SSL/SSH Inspection** profiles is all that can really be done in the policy.

The reason for having this inspection as part of the policy is the wide spread use of encryption by both legitimate and malicious actors. The legitimate users of the Internet use encryption to hide their information from snooping bad guy but the bad guys use encryption to hide their malicious content from being scanned for viruses and other malicious code by security devices.

By using the correct SSL certificates, the FortiGate can open up encrypted traffic and inspect it for malicious content that would otherwise make it past the other profiles because they couldn't read the encrypted traffic.

There are two basic types of inspection:

- Certificate inspection, which only looks at the certificate that encrypted the packets to make sure that it is a recognized and valid certificate.
- Full inspection, or deep inspection, that looks at all of the content of the packet. While more thorough, it also takes up more resources to perform.

HTTP Strict Transport Security (HSTS) Protocol

HSTS is a protocol used by Google and other web browsers to prevent man-in-the-middle attacks.

When performing deep inspection, the FortiGate intercepts the https traffic and would send its own self-signed CA certificate to the browser. If the browser is configured to use HSTS connections, it would refuse the FortiGate CA certificate since it is not on the trusted list for Google servers.

To keep the CA certificate from being refused, the HSTS settings should be cleared from the browser. Instructions for this vary between browsers.

SSH MITM deep inspection

Due to an increase, in recent years of vulnerabilities discovered in the SSH protocol, protections have been incorporated into FortiOS's Intrusion Prevention System (IPS) engine that will aid in protecting against malicious activity coming through the FortiGate against SSH access points. In addition to the protections offered by IPS, additional settings and functionality have been added to protect against man-in-the-middle (MITM) attacks that use SSH as the attack vector.

These protections include support for comprehensive security controls on SSH MITM deep inspections:

- SSH filter profile to control SSH tunnel types and filtering on SSH shell commands.
- SSH proxy policy that can apply a proxy firewall policy or firewall policy using SSH inspection, with user authentication to SSH sessions.
- Support for SSH tunnel policy access control for TCP/IP port forwarding traffic tunneled through SSH proxy. This allows IPS scanning to be applied to the tunneled traffic.
- Support the use of SSH trust status to detect and prevent SSH-based MITM attacks.

Support SSH proxy policy for SSH sessions

Policy check

To enable SSH proxy-policy or SSH tunnel-policy, `ssh-policy-check` or `ssh-tun-policy-check` must be enabled under `ssl-ssh-profile` for the corresponding firewall policy.

```
config firewall ssl-ssh-profile
  edit <name>
    config ssh
      set ssh-policy-check [disable|enable]
      set ssh-tun-policy-check [disable|enable]
    end
  end
```

SSH proxy option

Set the proxy type to `ssh` in `config firewall proxy-policy`.

```
config firewall proxy-policy
  edit <pol-id>
    set proxy ssh
    set action {accept|deny}
    set utm-status enable
    set ssh-filter-profile <profile_name>
  end
```

Authentication for SSH

When `user` or `user-group` is set in SSH proxy policy, firewall authentication can be implemented for SSH proxy traffic.

```
config authentication rule
edit <name>
set protocol ssh
end
```

Basic authentication scheme:

```
config authentication scheme
edit "ssh-active"
set method basic
set user-database "local" #or LDAP server
end
```



Under authentication for SSH, with basic authentication scheme, the client is expected to input username/password in the format of:

- `<firewall username>:<server username>`
- `<firewall password>:<server password>`

SSH-publickey authentication scheme:

```
config authentication scheme
edit "ssh-pkey"
set method ssh-publickey
set user-database "local"
set ssh-ca "server-ca"
end
```



The user name is embedded in `ssh-publickey`. The user group information will be retrieved if the public key is validated by CA.

FortiOS currently only supports validation by using CA. The CA needs to be configured under `config firewall ssh local-ca`.

Private-key is not required for client public-key authentication. Once the client offers the public-key signed by the set CA will be authenticated.

Both basic and SSH-publickey authentication scheme:

```
config authentication scheme
edit "ssh-pkey"
set method basic ssh-publickey
set user-database "local"
set ssh-ca "server-ca"
end
```

Support SSH tunnel policy to do access control for TCP/IP port forwarding traffic.

Add a proxy type `ssh-tunnel` into config firewall proxy-policy

```
config firewall proxy-policy
```

```
edit <pol-id>
  set proxy ssh-tunnel
  set action {accept | deny}
  set utm-status enable
```

The feature supports allowing or denying based on the IPS sensor/app-control applied to the traffic.

```
set ips-sensor <sensor_profile_name>
set application-list <application_control_list>
end
```

Support SSH trust to detect and prevent from SSH MITM attacks

Define trusted SSH host key for specific SSH server

```
config firewall ssh host-key
  edit <name>
    set status {trusted|revoked}
    set type {RSA|DSS|ECDSA}
    set nid <NID of ECDSA key>
    set ip <ip>
    set port <port>
    set hostname <name>
    set public-key <host key>
  end
```

Define trusted/untrusted CAs for host key signing.

Any host key signed by trust CA is trusted unless the host key is revoked.

```
config firewall ssh local-ca
  edit <name>
    set password <passwd>
    set public-key <public key>
    set private-key <private key>
    set source {built-in|user}
  end
```



- The system creates two built-in SSH CAs:
 - Fortinet_SSH_CA
 - Fortinet_SSH_CA_Untrusted.
- The CAs are used to re-sign a server host key with local host-key using trusted/untrusted CA when the server host key is trusted or untrusted.

Define local hostkey templates for trusted re-signing.

By default, the local hostkey templates are generated automatically.

```
config firewall ssh local-key
  edit <name>
    set status [trusted|revoked]
    set type [RSA|DSA|ECDSA|ED25519|RSA-CA|DSA-CA|ECDSA-CA|ED25519-CA]
    set nid <NID of ECDSA key>
    set ip <ip>
    set port <port>
```

```

set hostname <name>
set password <passwd>
set public-key <public key>
set private-key <private key>
set source {build-in|user}
end

```



Per VDOM SSH settings

```

config firewall ssh setting
set caname <trusted-ca>
set untrusted-caname <untrusted-ca>
set hostkey-rsa <hostkey-rsa>
set hostkey-dss <hostkey-dss>
set hostkey-ecdsa256 <hostkey-ecdsa256>
set hostkey-ecdsa384 <hostkey-ecdsa384>
set ed25519-key <ed25519-key>
set host-trusted-check {enable|disble}
end

```



- When a host key is trusted and signed by a CA, SSH proxy re-signs according to the type of hostkey using trusted CA.
- When a host is trusted but not signed, SSH proxy sends back according type of hostkey.
- When a host key is untrusted and signed by a CA, SSH proxy re-signs a temporary host key (1 hour) using untrusted CA.
- When a host is trusted but not signed, SSH proxy sends back a temporary host key (1 hour).

Add SSH filter profile table

Support options to block/log "x11-filter/ssh-shell/exec/port-forward/sftp"

```

config ssh-filter profile
edit <name>
set block {x11|shell|exec|port-forward|tun-forward|sftp|unknown}+
set log {x11|shell|exec|port-forward|tun-forward|sftp|unknown}+
end

```

Add shell command filters

```

config ssh-filter profile
edit <name>
config shell-commands
edit <id>
set type {simple|regex}
set pattern <cmd-string>
set action {block|allow}

```

```
set log {block|allow}
set alert {block|allow}
set severity {low|medium|high|critical}
end
set default-command-log {block|allow}
end
```

Mirroring SSL inspected traffic

It is possible to "mirror" or send a copy of traffic decrypted by SSL inspection to one or more FortiGate interfaces so that the traffic can be collected by a raw packet capture tool for archiving or analysis. The mirroring occurs after being processed by the SSL decoder and in the same point in the work flow as the decryption of application data. The decrypted application data is wrapped inside a TCP packet (with IP and Ethernet headers), and then sent to the mirror port.

This feature works when the inspection mode is set to flow-based or proxy-based, but not for explicit proxy.



Decryption, storage, inspection, and use decrypted content is subject to local privacy rules. Use of these features could enable malicious users with administrative access to your FortiGate to harvest sensitive information submitted using an encrypted channel.

In this example, the setting enables the policy to send all traffic decrypted by the policy to the FortiGate port1 and port2 interfaces.

```
config firewall policy
edit 0
set ssl-mirror enable
set ssl-mirror-intf port1 port2
end
```

Encryption strength for proxied SSH sessions

The level of SSH encryption can be set for SSH sessions on a per-profile basis.

Encryption Level	Description
compatible	This level allows for a broader set of encryption algorithms to be used and is better for compatibility.
high-encryption	This level will only allow AES-CTR, AES-GCM and high encryption algorithms to be used for the session.

Syntax:

```
config firewall ssl-ssh-profile
edit <profile name>
config ssh
set ssh-algorithm {compatible|high-encryption}
end
end
```

RPC over HTTP

How protocol options profiles and SSL inspection profiles handle RPC (Remote Procedure Calls) over HTTP traffic can be configured separately from normal HTTP traffic. The configuration is done in the CLI.

Configuration in protocol options

```
config firewall profile-protocol-options
edit 0
set rpc-over-http [disable|enable]
end
```

Configuration in SSL/SSH inspection

```
config firewall ssl-ssh-profile
edit deep inspection
set rpc-over-http [disable|enable]
end
```

NAT

NAT or Network Address Translation is the process that enables a single device such as a router or firewall to act as an agent between the Internet or Public Network and a local or private network. This “agent”, in real time, translates the source IP address of a device on one network interface, usually the Internal, to a different IP address as it leaves another interface, usually the interface connected to the ISP and the Internet. This enables a single public address to represent a significantly larger number of private addresses.

The origins of NAT

In order to understand NAT it helps to know why it was created. At one time, every computer that was part of a network had to have its own addresses so that the other computers could talk to it. There were a few protocols in use at the time, some of which were only for use on a single network, but of those that were routable, the one that had become the standard for the Internet was IP (Internet Protocol) version 4.

When IP version 4 addressing was created nobody had any idea how many addresses would be needed. The total address range was based on the concept of 2 to the 32nd power, which works out to be 4 294 967 296 potential addresses. Once you eliminate some of those for reserved addresses, broadcast addresses, network addresses, multicasting, etc., you end up with a workable scope of about 3.2 million addressees. This was thought to be more than enough at the time. The designers were not expecting the explosion of personal computing, the World Wide Web or smart phones. As of the beginning of 2012, some estimate the number of computers in the world in the neighborhood of 1 billion, and most of those computer users are going to want to be on the Internet or Search the World Wide Web. In short, we ran out of addresses.

This problem of an address shortage was realized before we actually ran out, and in the mid 1990s 2 technical papers called RFCs numbered 1631 (<http://www.ietf.org/rfc/rfc1631.txt>) and 1918 (<http://tools.ietf.org/html/rfc1918>), proposed components of a method that would be used as a solution until a new addressing methodology could be implemented across the Internet infrastructure. For more information on this you can look up IP version 6.

RFC 1631 described a process that would allow networking devices to translate a single public address to multiple private IP addresses and RFC 1918 laid out the use of the private addresses. The addresses that were on the Internet (Public IP addresses) could not be duplicated for them to work as unique addresses, but behind a

firewall, which most large institutions had, they could use their own Private IP addresses for internal use and the internal computers could share the external or Public IP address.

To give an idea on a small scale how this works, image that a company has a need for 200 computer addresses. Before Private IP addresses and NAT the company would have purchased a full Class C address range which would have been 254 usable IP addresses; wasting about 50 addresses. Now with NAT, that company only needs 1 IP address for its 200 computers and this leaves the rest of the IP addresses in that range available for other companies to do the same thing.

NAT gives better value than it would first appear because it is not 253 companies that can use 254 addresses but each of those 254 companies could set up their networking infrastructures to use up to thousands of Private IP addresses, more if they don't all have to talk to the Internet at the same time. This process enabled the Internet to keep growing even though we technically have many more computers networked than we have addresses.

Dynamic NAT

Dynamic NAT maps the private IP addresses to the first available Public Address from a pool of possible Addresses. In the FortiGate firewall this can be done by using IP Pools.

Overloading

This is a form of Dynamic NAT that maps multiple private IP address to a single Public IP address but differentiates them by using a different port assignment. This is probably the most widely used version of NAT. This is also referred to as PAT (Port Address Translation) or Masquerading.

An example would be if you had a single IP address assigned to you by your ISP but had 50 or 60 computers on your local network.

Say the internal address of the interface connected to the ISP was 256.16.32.65 (again an impossible address) with 256.16.32.64 being the remote gateway. If you are using this form of NAT any time one of your computers accesses the Internet it will be seen from the Internet as 256.16.32.65. If you wish to test this go to 2 different computers and verify that they each have a different private IP address then go to a site that tells you your IP address such as www.ipchicken.com. You will see that the site gives the same result of 256.16.32.65, if it existed, as the public address for both computers.

As mentioned before this is sometimes called Port Address Translation because network device uses TCP ports to determine which internal IP address is associated with each session through the network device. For example, if you have a network with internal addresses ranging from 192.168.1.1 to 192.168.1.255 and you have 5 computers all trying to connect to a web site which is normally listening on port 80 all of them will appear to the remote web site to have the IP address of 256.16.32.65 but they will each have a different sending TCP port, with the port numbers being somewhere between 1 and 65 535, although the port numbers between 1 to 1024 are usually reserved or already in use. So it could be something like the following:

192.168.1.10	256.16.32.65:	port 486
192.168.1.23	256.16.32.65:	port 2409
192.168.1.56	256.16.32.65:	port 53763
192.168.1.109	256.16.32.65:	port 5548
192.168.1.201	256.16.32.65:	port 4396

And the remote web server would send the responding traffic back based on those port numbers so the network device would be able to sort through the incoming traffic and pass it on to the correct computer.

Overlapping

Because everybody is using the relative same small selection of Private IP addresses it is inevitable that there will be two networks that share the same network range that will need to talk with each other. This happens most often over Virtual Private Networks or when one organization ends up merging with another. This is a case where a private IP address may be translated into a different private IP address so there are no issues with conflict of addresses or confusion in terms of routing.

An example of this would be when you have a Main office that is using an IP range of 172.16.0.1 to 172.20.255.255 connecting through a VPN to a recently acquired branch office that is already running with an IP range of 172.17.1.1 to 172.17.255.255. Both of these ranges are perfectly valid but because the Branch office range is included in the Main Office range any time the system from the Main office try to connect to an address in the Branch Office the routing the system will not send the packet to the default gateway because according to the routing table the address is in its own subnet.

The plan here would be to NAT in both directions so that traffic from neither side of the firewall would be in conflict and they would be able to route the traffic. Everything coming from the Branch Office could be assigned an address in the 192.168.1.1 to 192.168.1.255 range and everything from the Main office going to the Branch Office could be assigned to an address in the 192.168.10.1 to 192.168.10.255 range.

Static NAT

In Static NAT one internal IP address is always mapped to the same public IP address.

In FortiGate firewall configurations this is most commonly done with the use of Virtual IP addressing.

An example would be if you had a small range of IP addresses assigned to you by your ISP and you wished to use one of those IP address exclusively for a particular server such as an email server.

Say the internal address of the Email server was 192.168.12.25 and the Public IP address from your assigned addresses range from 256.16.32.65 to 256.16.32.127. Many readers will notice that because one of the numbers is above 255 that this is not a real Public IP address. The Address that you have assigned to the interface connected to your ISP is 256.16.32.66, with 256.16.32.65 being the remote gateway. You wish to use the address of 256.16.32.70 exclusively for your email server.

When using a Virtual IP address you set the external IP address of 256.16.32.70 to map to 192.168.12.25. This means that any traffic being sent to the public address of 256.16.32.70 will be directed to the internal computer at the address of 192.168.12.25

When using a Virtual IP address, this will have the added function that when ever traffic goes from 192.168.12.25 to the Internet it will appear to the recipient of that traffic at the other end as coming from 256.16.32.70.

You should note that if you use Virtual IP addressing with the Port Forwarding enabled you do not get this reciprocal effect and must use IP pools to make sure that the outbound traffic uses the specified IP address.

Benefits of NAT

More IP addresses available while conserving public IP addresses

As explained earlier, this was the original intent of the technology and does not need to be gone into further.

Financial savings

Because an organization does not have to purchase IP addresses for every computer in use there is a significant cost savings due to using the process of Network Address Translation.

Security enhancements

One of the side benefits of the process of NAT is an improvement in security. Individual computers are harder to target from the outside and if port forwarding is being used computers on the inside of a firewall are less likely to have unmonitored open ports accessible from the Internet.

Ease of compartmentalization of your network

With a large available pool of IP addresses to use internally a network administrator can arrange things to be compartmentalized in a rational and easily remembered fashion and networks can be broken apart easily to isolate for reasons of network performance and security.

Example

You have a large organization that for security reasons has certain departments that do not share network resources.

You can have the main section of the organization set up as follows;

Network Devices	192.168.1.1 to 192.168.1.25
Internal Servers	192.168.1.26 to 192.168.1.50
Printers	192.168.1.51 to 192.168.1.75
Administration Personnel	192.168.1.76 to 192.168.1.100
Sales People	192.168.1.101 to 192.168.1.200
Marketing	192.168.1.201 to 192.168.1.250

You could then have the following groups broken off into separate subnets:

Accounting	192.168.100.1 to 192.168.100.255
Research and Development	172.16.1.1 to 172.16.255.255
Executive Management	192.168.50.1 to 192.168.50.255
Web sites and Email Servers	10.0.50.1 to 10.0.50.255

These addresses do not have to be assigned right away but can be used as planned ranges.

NAT in transparent mode

Similar to operating in NAT mode, when operating a FortiGate unit in transparent mode you can add security policies and:

- Enable NAT to translate the source addresses of packets as they pass through the FortiGate unit.
- Add virtual IPs to translate destination addresses of packets as they pass through the FortiGate unit.
- Add IP pools as required for source address translation

A FortiGate unit operating in transparent mode normally has only one IP address - the management IP. To support NAT in transparent mode, you can add a second management IP. These two management IPs must be on different subnets. When you add two management IP addresses, all FortiGate unit network interfaces will respond to connections to both of these IP addresses.

Use the following steps to configure NAT in transparent mode:

1. Add two management IPs
2. Add an IP pool to the WAN1 interface
3. Add an Internal to WAN1 security policy

You can add the security policy from the web-based manager and then use the CLI to enable NAT and add the IP pool.

The usual practice of NATing in transparent mode makes use of two management IP addresses that are on different subnets, but this is not an essential requirement in every case.

If there is a router between the client systems and the FortiGate unit you can use the router's capabilities of tracking sessions to assign NATed addresses from an IP pool to the clients even if the assigned address don't belong to a subnet on your network.

Example

Client computer has an IP address of 1.1.1.33 on the subnet 1.1.1.0/24.

Router "A" sits between the client computer and the FortiGate (in transparent mode) with the IP address of 1.1.1.1 on the client's side of the router and the IP address of 192.168.1.211 on the FortiGate's side of the router.

Use NAT to assign addresses from an address pool of 9.9.9.1 to 9.9.9.99 to traffic coming from gateway of 192.168.1.211.

To enable the return traffic to get to the original computer, set up a static route that assigns any traffic with a destination of 9.9.9.0/24 to go through the 192.168.1.211 gateway. As long as the session for the outgoing traffic has been maintained, communication between the client computer and the external system on the other side of the FortiGate will work.

Central NAT table

The central NAT table enables you to define, and control with more granularity, the address translation performed by the FortiGate unit. With the NAT table, you can define the rules which dictate the source address or address group and which IP pool the destination address uses.

While similar in functionality to IP pools, where a single address is translated to an alternate address from a range of IP addresses, with IP pools there is no control over the translated port. When using the IP pool for source NAT, you can define a fixed port to guarantee the source port number is unchanged. If no fix port is defined, the port translation is randomly chosen by the FortiGate unit. With the central NAT table, you have full control over both the IP address and port translation.

The FortiGate unit reads the NAT rules in a top-down methodology, until it hits a matching rule for the incoming address. This enables you to create multiple NAT policies that dictate which IP pool is used based on the source

address. The NAT policies can be rearranged within the policy list as well. NAT policies are applied to network traffic after a security policy.

NAT64 and NAT46

NAT64 and NAT46 are the terms used to refer to the mechanism that allows IPv6 addressed hosts to communicate with IPv4 addressed hosts and vice versa. Without such a mechanism an IPv6 node on a network such as a corporate LAN would not be able to communicate with a web site that was still in a IPv4 only environment and IPv4 environments would not be able to connect to IPv6 networks.

One of these setups involves having at least 2 interfaces, 1 on an IPv4 network and 1 on an IPv6 network. The NAT64 server synthesizes AAAA records, used by IPv6 from A records used by IPv4. This way client-server and peer to peer communications will be able to work between an IPv6 only client and an IPv4 server without making changes to either of the end nodes in the communication transaction. The IPv6 network attached to the FortiGate unit should be a 32 bit segment, (for instance 64:ff9b::/96, see RFC 6052 and RFC 6146). IPv4 address will be embedded into the communications from the IPv6 client.

Because the IPv6 range of addresses is so much larger than the IPv4 range, a one to one mapping is not feasible. Therefore the NAT64 function is required to maintain any IPv6 to IPv4 mappings that it synthesizes. This can be done either statically by the administrator or automatically by the service as the packets from the IPv6 network go through the device. The first method would be a stateless translation and the second would be a stateful translation. NAT64 is designed for communication initiated from IPv6 hosts to IPv4 addresses. It is address mapping like this that allows the reverse to occur between established connections. The stateless or manual method is an appropriate solution when the NAT64 translation is taking place in front of legacy IPv4 servers to allow those specific servers to be accessed by remote IPv6-only clients. The stateful or automatic solution is best used closer to the client side when you have to allow some specific IPv6 clients to talk to any of the IPv4-only servers on the Internet.

There are currently issues with NAT64 not being able to make everything accessible. Examples would be SIP, Skype, MSN, Goggle talk, and sites with IPv4 literals. IPv4 literals being IPv4 addresses that are imbedded into content rather than a FQDN.

Policies that employ NAT64 or NAT46 can be configured from the web-based manager as long as the feature is enabled using the Features setting found at **System > Config > Features**.

- To create a NAT64 policy go to **Policy & Objects > NAT64 Policy** and select **Create New**.
- To create a NAT46 policy go to **Policy > NAT46 Policy** and select **Create New**.

The difference between these NAT policies and regular policies is that there is no option to use the security profiles and sensors.



NAT64 CLAT traffic is now supported by the FortiGate. CLAT traffic comes from devices that use the SIIT translator that plays a part in affecting IPv6 - IPv4 NAT translation.

NAT64 CLAT

NAT64 CLAT traffic is supported by FortiOS. CLAT traffic comes from devices that use the SIIT translator that plays a part in affecting IPv6 - IPv4 NAT translation.

NAT66

NAT66 is Network Address Translation between 2 IPv6 network. The basic idea behind NAT66 is no different than the regular NAT between IPv4 networks that we are all used to. The difference are in the mechanics of how it is performed, mainly because of the complexity and size of the addresses that are being dealt with.

In an IPv4 world, the reason for the use of NAT was usually one or a combination of the following 3 reasons:

- Improved security - actual addresses behind NAT are virtually hidden
- Amplification of addresses - hundreds of computers can use as little as a single public IP address
- Internal address stability - there is control of internal addressing. The addresses can stay the same even if Internet Service Providers change.

In these days of security awareness the protective properties of NAT are not something that are not normally depended on by themselves to defend a network and with the vastly enlarged IPv6 address scope there is no longer a need to amplify the available addresses. However, the desire to have internal address control still exists. The most common reason for using NAT66 is likely to be the maintaining of the existing address scheme of the internal network despite changes outside of it. Imagine that you have an internal network of 2000 IP addresses and one day the company changes its ISP and thus the addresses assigned to it. Even if most of the addressing is handled by DHCP, changing the address scheme is going to have an impact on operations.

Addressing stability can be achieved by:

- Keeping the same provider - this would depend on the reason for the change. If the cost of this provider has become too expensive this is unlikely. If the ISP is out of business it becomes impossible.
- Transfer the addresses from the old provider to the new one - There is little motivation for an ISP to do you a favor for not doing business with them.
- Get your own autonomous system number - this can be too expensive for smaller organizations.
- NAT - this is the only one on the list that is in the control of IT.

There are differences between NAT66 and IPv4 NAT. Because there is no shortage of addresses most organizations will be given a /48 network that can be translated into another /48 network. This allows for a one to one translation, no need for port forwarding. This is a good thing because port forwarding is more complicated in IPv6. In fact, NAT66 will actually just be the rewriting of the prefix on the address.

Example

If your current IPv6 address is

```
2001:db8:cafe::/48
```

you could change it to

```
2001:db8:fea7::/48
```

There is an exception to the one to one translation. NAT66 cannot translate internal networks that contain 0xffff in bits 49 through 63 - this is due to the way checksums are calculated in TCP/IP: they use the one's-complement representation of numbers which assigns the value zero to both 0x0000 and 0xffff.

How FortiOS differentiates sessions when NATing

The basics of NAT are fairly simple. Many private addresses get translated into a smaller number of public addresses, often just one. The trick is how the FortiGate keeps track of the return traffic because the web server, or what ever device that was out on the Internet is going to be sending traffic back not to the private address behind the FortiGate but to the IP address of the interface on the public side of the FortiGate.

The way this is done is by making each session unique. Most of the attributes that are available in the network packets cannot be changed without changing where the packet will go but because the source port has to be changed anyway in case two computer on the network used the same source port this is a useful way of making each listing of network attributes a unique combination. As a packet goes through the NAT process FortiOS assigns different source ports for each of the internally initiated sessions and keeping track of which port was used for each device in a database until the session has ended. It then becomes a matter of how the port number is selected.

In a very simple example of an environment using NAT, we will use a fictitious university with a rather large student population. So large in fact that they use a subnet of `10.0.0.0/8` as their subnet for workstation IP addresses. All of these private IP addresses are NATed out a single IP address. To keep the number of numeric values in this example from getting to a confusing level, we'll just use "u.u.u.1" to refer to the public IP address of the University and the IP address of the web server on the Internet will be "w.w.w.1".

Student A (IP address `10.1.1.56`) sends an HTML request to a web server on the Internet with the IP address `w.w.w.1`. The applicable networking information in the packet breaks down as follows:

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	<code>10.1.1.56</code>	<code>u.u.u.1</code>
Destination IP address or dst-ip:	<code>w.w.w.1</code>	<code>w.w.w.1</code>
Source port or src-port:	10000	46372
Destination port or dst-port	80	80

The source IP address is now that of the public facing interface of the FortiGate and source port number is an unused TCP port number on the FortiGate chosen by the FortiGate. Of these variable the only one the that FortiGate can really change and still have the packet reach the correct destination, in both directions, is the source port number.

There are a few methods of assigning the port number. First we'll look at the methods that are or have been used in the industry but aren't used by Fortinet.

Global pool

This method of differentiation focuses on the attribute of the source port number. In this approach a single pool of potential port numbers is set aside for the purposes of NAT. As a pool number is assigned, it is removed from the pool so that two sessions from different computers can not using the same port number. Once the session is over and no longer in use by the computer, the port number is put back into the pool where it can be assigned again.

Example global pool:

	Hexidecimal	Decimal
Start or range	0x7000	28672

	Hexidecimal	Decimal
End end of range	0xF000	61440
Possible ports in range	215	32768

This is a simple approach to implement and is good if the number of connections is unlikely to reach the pool size. It would be okay for home use, but our example is for a university using 10.1.1.0/8 as a subnet. That means 16,777,214 possible IP addresses; more than this method can handle.

Fortinet does not use this method.

Global per protocol

This method uses the attributes source port number and type of protocol to differentiate between sessions. This approach is a variation of the first one. An additional piece of information is referred to in the packet that describes the protocol. For instance UDP or TCP. This could effectively double the number of potential addresses to NAT.

Example:

Here are two possible packets that would be considered different by the FortiGate so that any responses from the web server would make it back to their correct original sender.

From Student A

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.1.1.56	u.u.u.1
Destination IP address or dst-ip:	w.w.w.1	w.w.w.1
Protocol	tcp	tcp
Source port or src-port:	10000	46372
Destination port or dst-port	80	80

From Student B

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.5.1.233	u.u.u.1
Destination IP address or dst-ip:	w.w.w.1	w.w.w.1
Protocol	udp	udp
Source port or src-port:	26785	46372
Destination port or dst-port	80	80

Even though the source port is the same, because the protocol is different they are considered to be from different sessions and different computers.

The drawback is that it would depend on the protocols being used be evenly distributed between TCP and UDP. Even if this was the case the number would only double; reaching an upper limit of 65,536 possible connections. That number is still far short of the possible more than 16 million for an IP subnet with an eight bit subnet mask like the one in our example.

Fortinet does not use this method.

Per NAT IP pool

This approach adds on to the previous one by adding another variable. In this case that variable is the IP addresses on the public side of the FortiGate. By having a pool of IP addresses to assign as the source IP address when NATing, the same number that was potentially available for the Global per protocol method can be multiplied by the number of external IP addresses in the pool. If you can assign a second IP address to the pool, you can double the potential number of sessions.

Example:

In this example it will be assumed that the FortiGate has 2 IP addresses that it can use. This could happen either by using two ISPs, or by having a pool of IP addresses assigned to a single interface. For simplicity will refer to these IP public IP addresses as `u.u.u.1` and `u.u.u.2`.

Here are two possible packets that would be considered different by the FortiGate so that any responses from the web server would make it back to their correct original sender.

From Student A

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.1.1.56	u.u.u.1
Destination IP address or dst-ip:	w.w.w.1	w.w.w.1
Protocol	tcp	tcp
Source port or src-port:	10000	46372
Destination port or dst-port	80	80

From Student B

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.5.1.233	u.u.u.2
Destination IP address or dst-ip:	w.w.w.1	w.w.w.1
Protocol	tcp	tcp

Attribute	Original Packet	Packet after NATing
Source port or src-port:	26785	46372
Destination port or dst-port	80	80

In this example we even made the protocol the same. After the NATing process all of the variables are the same except the source address. This is still going to make it back to the original sender.

The drawback is that if you have only one IP address for the purposes of NATing this method does not gain you anything over the last method. Or if you do have multiple IP addresses to use it will still take quite a few to reach the 16 million possible that the subnet is capable of handling.

Fortinet does not use this method.

Per NAT IP, destination IP, port, and protocol

This is the approach that FortiOS uses.

It uses all of the differentiation point of the previous methods, NAT IP, port number and protocol, but the additional information point of the destination IP is also used. So now the network information points in the packet that the FortiGate keeps in its database to differentiate between sessions is:

- Public IP address of the FortiGate assigned by NATing
- Protocol of the traffic
- Source port assigned by the FortiGate
- Destination IP address of the packet

The last one is an especially good way to differentiate because as a theoretical number, the upper limit on that is the numbers of Public IP addresses on the whole of the Internet. Chances are that while a large number of session from inside the University will be going to a small group of sites such as Google, Youtube, Facebook and some others it is unlikely that they will all be going to them at the same time.

Example:

In this example it will be assumed that the FortiGate has only one IP address. Two possible packets will be described. The only difference in the attributes recorded will be the destination of the HTML request. These packets are still considered to be from different sessions and any responses will make it back to the correct computer.

From Student A

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.1.1.56	u.u.u.1
Destination IP address or dst-ip:	w.w.w.1	w.w.w.1
Protocol	tcp	tcp

Attribute	Original Packet	Packet after NATing
Source port or src-port:	10000	46372
Destination port or dst-port	80	80

From Student B

Attribute	Original Packet	Packet after NATing
Source IP address or src-ip	10.5.1.233	u.u.u.1
Destination IP address or dst-ip:	w.w.w.2	w.w.w.2
Protocol	tcp	tcp
Source port or src-port:	26785	46372
Destination port or dst-port	80	80

The reason that these attributes are used to determine differentiation between traffic is based on how the indexes for the sessions are recorded in the database. When a TCP connection is made through a FortiGate unit, a session is created and two indexes are created for the session. The FortiGate unit uses these indexes to guide matching traffic to the session.

This following could be the session record for the TCP connection in the first example.

Attribute	Outgoing Traffic	Returning Traffic
Source IP address	10.78.33.97 (internal address)	w.w.w.1
Destination address	w.w.w.1	u.u.u.1
Protocol	tcp	tcp
Source port	10000 (from original computer) 46372 (assigned by NAT)	80
Destination port	80	46372 (FortiGate assigned port)

Using the FortiGate's approach for session differentiation, FortiOS only has to ensure that the assigned port, along with the other four attributes is a unique combination to identify the session. So for example, if Student A simultaneously makes a HTTP(port 80) connection and a HTTPS(port 443) connection the same web server this would create another session and the index in the reply direction would be:

Attribute	Outgoing Traffic	Returning Traffic
Source IP address	10.78.33.97 (internal address)	w.w.w.1
Destination address	w.w.w.1	u.u.u.1
Protocol	tcp	tcp
Source port	10000 (from original computer) 46372 (assigned by NAT)	443
Destination port	443	46372 (FortiGate assigned port)

These two sessions are different and acceptable because of the different source port numbers on the returning traffic or the destination port depending on the direction of the traffic.

Calculations for possible session numbers

The result of using these four attributes instead of just the one that was originally used is a large increase in the number of possible unique combinations. For those who love math, the maximum number of simultaneous connections that can be supported is:

$$N \times R \times P \times D \times Dp$$

where:

- **N** is the number of NAT IP addresses
- **R** is the port range,
- **P** is the number of protocols,
- **D** is the number of unique destination IP addresses
- **Dp** the number of unique destination ports.

As a rough example let's do some basic calculations

- **N** - In our existing example we have already stated that there is only one public IP address that is being used by NAT. Realistically, for a university this number would likely be larger, but we're keeping it simple.

$$N = 1$$

R - The port range for our example has already been describe and we will keep it the same.

$$R = 32768$$

P - While there are a few protocols that are involved in Internet traffic we will limit this calculation just to TCP traffic.

$$P = 1$$

D - As mentioned before the number of unique destination addresses is growing larger every day, so figuring out the upper limit of that number would be difficult to say the least. Instead we will make the assumption that most of the university students, do to their shared interest and similar demographic will concentrate most of their web browsing to the same sites; sites such as YouTube, Facebook, Google, Twitter, Instagram, Wikipedia etc. This is

not even taking into account the fact that many of these popular sites use load balancing and multiple IP addresses. As an arbitrary number let's use the number 25.

D = 25

Dp - To keep things simple it is tempting to limit the destination port to port 80, the one that many associate with web browsing, but this would not be realistic. the use of HTTPS, port 443 is on the rise. There is also email, DNS, FTP, NTP and a number of other background services that we use without thinking too closely about. Let's keep it small and say ten of them.

Dp = 10

The math on this very conservative calculation is:

$1 \times 32768 \times 1 \times 25 \times 10 = 8,192,000$ possible NAT sessions

When you take into account that the chances of everybody being online at the same time, going only to one of those 25 sites and not millions of others, and using only TCP not UDP or any of the other protocols, it starts to look like this method may provide enough potential unique sessions even for a subnet as large as the one described.

IP pools

IP Pools are a mechanism that allow sessions leaving the FortiGate Firewall to use NAT. An IP pool defines a single IP address or a range of IP addresses to be used as the source address for the duration of the session. These assigned addresses will be used instead of the IP address assigned to that FortiGate interface.



When using IP pools for NATing, there is a limitation that must be taken into account. In order for communication to be successful in both directions, it is normal for the source address in the packet header assigned by the NAT process to be an address that is associated with the interface that the traffic is going through. For example, if traffic is going out an interface with the IP address 172.16.100.1, packets would be NATed so that the source IP address would be 172.16.100.1. This way the returning traffic will be directed to the same interface on the same FortiGate that the traffic left from. Even if the packets are assigned a source address that is associated with another interface on the same FortiGate this can cause issues with asymmetrical routing. It is possible to configure the NATed source IP address to be different than the IP address of the interface but you have to make sure that the routing rules of the surrounding network devices take this unorthodox approach into consideration.

There are 4 types of IP Pools that can be configured on the FortiGate firewall:

- One-to-One - in this case the only internal address used by the external address is the internal address that it is mapped to.
- Overload - this is the default setting. Internal addresses other than the one designated in the policy can use this address for the purposes of NAT.
- Fixed Port Range - rather than a single address to be used, there is a range of addresses that can be used as the NAT address. These addresses are randomly assigned as the connections are made.
- Port Block Allocation - this setting is used to allocate a block of port numbers for IP pool users. Two variables will also have to be set. The block size can be set from 64 to 4096 and as the name implies describes the number of ports in one block of port numbers. The number of blocks per user determines how many of these blocks will be assigned. This number can range from 1 to 128.



Be careful when calculating the values of the variables. The maximum number of ports that are available on an address is 65,536. If you chose the maximum value for both variables you will get a number far in excess of the available port numbers.

$$4096 \times 128 = 524,288$$

One of the more common examples is when you have an email server behind your FortiGate firewall and the range of IP addresses assigned to you by your ISP is more than one. If an organization is assigned multiple IP addresses it is normally considered a best practice to assign a specific address other than the one used for the Firewall to the mail server. However, when normal NAT is used the address assigned to the firewall is also assigned to any outbound sessions. Anti-spam services match the source IP address of mail traffic that they receive to the MX record on DNS servers as an indicator for spam. If there is a mismatch the mail may not get through so there is a need to make sure that the NATed address assigned matches the MX record.

You can also use the Central NAT table as a way to configure IP pools.

Source IP address and IP pool address matching when using a range

When the source addresses are translated to an IP pool that is a range of addresses, one of the following three cases may occur:

Scenario 1:

The number of source addresses equals that of IP pool addresses

In this case, the FortiGate unit always matches the IP addressed one to one.

If you enable fixed port in such a case, the FortiGate unit preserves the original source port. This may cause conflicts if more than one security policy uses the same IP pool, or the same IP addresses are used in more than one IP pool.

Scenario 2:

The number of source addresses is more than that of IP pool addresses

In this case, the FortiGate unit translates IP addresses using a wrap-around mechanism. If you enable fixed port in such a case, the FortiGate unit preserves the original source port. But conflicts may occur since users may have different sessions using the same TCP 5 tuples.

Scenario 3:

The number of source addresses is fewer than that of IP pool addresses

In this case, some of the IP pool addresses are used and the rest of them are not be used.

ARP replies

If a FortiGate firewall interface IP address overlaps with one or more IP pool address ranges, the interface responds to ARP requests for all of the IP addresses in the overlapping IP pools. For example, consider a FortiGate unit with the following IP addresses for the port1 and port2 interfaces:

- port1 IP address: 1.1.1.1/255.255.255.0 (range is 1.1.1.0-1.1.1.255)
- port2 IP address: 2.2.2.2/255.255.255.0 (range is 2.2.2.0-2.2.2.255)

And the following IP pools:

- IP_pool_1: 1.1.1.10-1.1.1.20
- IP_pool_2: 2.2.2.10-2.2.2.20
- IP_pool_3: 2.2.2.30-2.2.2.40

The port1 interface overlap IP range with IP_pool_1 is:

$(1.1.1.0-1.1.1.255) \text{ and } (1.1.1.10-1.1.1.20) = 1.1.1.10-1.1.1.20$

The port2 interface overlap IP range with IP_pool_2 is:

$(2.2.2.0-2.2.2.255) \text{ \& } (2.2.2.10-2.2.2.20) = 2.2.2.10-2.2.2.20$

The port2 interface overlap IP range with IP_pool_3 is:

$(2.2.2.0-2.2.2.255) \text{ \& } (2.2.2.30-2.2.2.40) = 2.2.2.30-2.2.2.40$

And the result is:

- The port1 interface answers ARP requests for 1.1.1.10-1.1.1.20
- The port2 interface answers ARP requests for 2.2.2.10-2.2.2.20 and for 2.2.2.30-2.2.2.40

Select Enable NAT in a security policy and then select Dynamic IP Pool. Select an IP pool to translate the source address of packets leaving the FortiGate unit to an address randomly selected from the IP pool. Whether or not the external address of an IP Pool will respond to an ARP request can be disabled. You might want to disable the ability to respond to ARP requests so that these address cannot be used as a way into your network or show up on a port scan.

IP pools and zones

Because IP pools are associated with individual interfaces IP pools cannot be set up for a zone. IP pools are connected to individual interfaces.

Fixed port

Some network configurations do not operate correctly if a NAT policy translates the source port of packets used by the connection. NAT translates source ports to keep track of connections for a particular service.

However, enabling the use of a fixed port means that only one connection can be supported through the firewall for this service. To be able to support multiple connections, add an IP pool, and then select Dynamic IP pool in the policy. The firewall randomly selects an IP address from the IP pool and assigns it to each connection. In this case, the number of connections that the firewall can support is limited by the number of IP addresses in the IP pool.

Match-VIP

The match-vip feature allows the FortiGate unit to log virtual IP traffic that gets implicitly dropped. This feature eliminates the need to create two policies for virtual IPs; one that allows the virtual IP, and the other to get proper log entry for DROP rules.

For example, you have a virtual IP security policy and enabled the match-vip feature; the virtual IP traffic that is not matched by the policy is now caught.

The match-vip feature is available only in the CLI. By default, the feature is disabled.

Services and TCP ports

There are a number of different services and protocols in use on the Internet. The most commonly known is HTTP which is used by web servers to transmit requests and responses for unencrypted web pages. These services are set up to listen for requests on a numbered port. These services and protocols can use any port from 1 to 65,535. To keep things simple for everyone a large number of the more commonly used services started using a standardized list of ports. For instance, though it is not required, by default, most web servers listen for HTTP requests on port 80 and by default, web browsers will send HTTP traffic to port 80. If you wish to use another port such as 8080 you would put “:8080” at the end of the URL to indicate that you want the browser to use 8080 instead of the default port.

Example

Default URL for HTTP traffic when the web server is listening on the standard HTTP port:

`http://fortinet.com`

URL to the same address when the web server is listening for HTTP traffic on port 8080

`http://fortinet.com:8080`

Services represent typical traffic types and application packets that pass through the FortiGate unit. Firewall services define one or more protocols and port numbers associated with each service. Security policies use service definitions to match session types. You can organize related services into service groups to simplify your security policy list.

Many well-known traffic types have been predefined on the FortiGate unit. If there is a service that does not appear on the list you can create a service or edit an existing one. You need to know the ports, IP addresses or protocols of that particular service or application uses, to create a service.

Best Practices



While you can edit a predefined service it is best to leave those ones alone and create a new service and name it something similar such as the same service name with a descriptive identifier appended.

Based on the previous example, instead of the name “HTTP” you could name the service “HTTP8080” or use the application that is using that port, “HTTP-Application”.

Protocol types

One of the fundamental aspects of a service is the type of protocol that use used to define it. When a service is defined one of the following categories of protocol needs to be determined:

- TCP/UDP/SCTP
- ICMP
- ICMPv6
- IP

Depending on which of these protocol categories is choose another set of specifications will can also be defined.

Protocol Type	Related specifications
TCP/UDP/SCTP	This is the most commonly used service protocol category. Once this category has been selected the other available options to choose are an address, either IP or FQDN, and the protocol and port number. The protocol will be TCP, UDP or SCTP.
ICMP or ICMP6	When ICMP or ICMP6 is chosen the available options are the ICMP Type and its code.
IP	When IP is the chosen protocol type the addition option is the Protocol Number.

TCP/UDP/SCTP

TCP

Transmission Control Protocol (TCP) is one of the core or fundamental protocols of the Internet. It is part of the Transport Layer of the OSI Model. It is designed to provide reliable delivery of data from a program on one device on the network or Internet to another program on another device on the network or Internet. TCP achieves its reliability because it is a connection based protocol. TCP is stream-oriented. It transports streams of data reliably and in order.

TCP establishes a prior connection link between the hosts before sending data. This is often referred to as the handshake. Once the link is established the protocol uses checks to verify that the data transmitted. If an error check fails the data is retransmitted. This makes sure that the data is getting to the destination error free and in the correct order so that it can be put back together into a form that is identical to the way they were sent.

TCP is configured more for reliability than for speed and because of this TCP will likely be slower than a connectionless protocol such as UDP. This is why TCP is generally not used for real time applications such as voice communication or online gaming.

Some of the applications that use TCP are:

- World Wide Web (HTTP and HTTPS)
- Email (SMTP, POP3, IMAP4)
- Remote administration (RDP)
- File transfer (FTP)

UDP

User Datagram Protocol (UDP) like TCP is one of the core protocols of the Internet and part of the Transport Layer of the OSI Model. UDP is designed more for speed than reliability and is generally used for different applications than TCP. UDP sends messages, referred to as datagrams across the network or Internet to other hosts without establishing a prior communication link. In other words, there is no handshake.

UDP is an unreliable service as the datagrams can arrive out of order, duplicated or go missing without any mechanism to verify them. UDP works on the assumption that any error checking is done by the application or is not necessary for the function of the application. This way it avoids the overhead that is required to verify the integrity of the data.

This lack of overhead improves the speed of the data transfer and is why UDP is often used by applications that are time sensitive in nature. UDP's stateless nature is also great for applications that answer a large number of small queries from a large number of clients.

Common uses for UDP are:

- Domain Name Resolution (DNS)
- Time (NTP)
- Streaming media (RTSP, RTP and RTCP)
- Telephone of the Internet (VoIP)
- File Transfer (TFTP)
- Logging (SNMP)
- Online games (GTP and OGP)

SCTP

Stream Control Transmission Protocol (SCTP) is part of the Transport Layer of the OSI Model just like TCP and UDP and provides some of the features of both of those protocols. It is message or datagram orientated like UDP but it also ensures reliable sequential transport of data with congestion control like TCP.

SCTP provides the following services:

- Acknowledged error-free non-duplicated transfer of user data
- Data fragmentation to conform to discovered path MTU size
- Sequenced delivery of user messages within multiple streams, with an option for order-of-arrival delivery of individual user messages
- Optional bundling of multiple user messages into a single SCTP packet
- Network-level fault tolerance through supporting of multi-homing at either or both ends of an association
- Congestion avoidance behavior and resistance to flooding and masquerade attacks

SCTP uses multi-streaming to transport its messages which means that there can be several independent streams of messages traveling in parallel between the points of the transmission. The data is sent out in larger chunks of data than is used by TCP just like UDP but the messages include a sequence number within each message in the same way that TCP does so that the data can be reassembled at the other end of the transmission in the correct sequence without the data having to arrive in the correct sequence.

SCTP is effective as the transport protocol for applications that require monitoring and session-loss detection. For such applications, the SCTP path and session failure detection mechanisms actively monitor the connectivity of the session. SCTP differs from TCP in having multi-homing capabilities at either or both ends and several streams within a connection, typically referred to as an association. A TCP stream represents a sequence of bytes; an SCTP stream represents a sequence of messages.

Some common applications of SCTP include supporting transmission of the following protocols over IP networks:

- SCTP is important in 3G and 4G/LTE networks (for example, HomeNodeB = FemtoCells)
- SS7 over IP (for example, for 3G mobile networks)
- SCTP is also defined and used for SIP over SCTP and H.248 over SCTP
- Transport of Public Switched Telephone Network (PSTN) signaling messages over IP networks.

SCTP is a much newer protocol. It was defined by the IETF Signaling Transport (SIGTRAN) working group in 2000. It was introduced by RFC 3286 and more fully define by RFC 4960.

The FortiGate firewall can apply security policies to SCTP sessions in the same way as TCP and UDP sessions. You can create security policies that accept or deny SCTP traffic by setting the service to “ALL”. FortiOS does not include pre-defined SCTP services. To configure security policies for traffic with specific SCTP source or destination ports you must create custom firewall services for SCTP.

FortiGate units route SCTP traffic in the same way as TCP and UDP traffic. You can configure policy routes specifically for routing SCTP traffic by setting the protocol number to 132. SCTP policy routes can route SCTP traffic according to the destination port of the traffic if you add a port range to the policy route.

You can configure a FortiGate unit to perform stateful inspection of different types of SCTP traffic by creating custom SCTP services and defining the port numbers or port ranges used by those services. FortiGate units support SCTP over IPv4. The FortiGate unit performs the following checks on SCTP packets:

- Source and Destination Port and Verification Tag.
- Chunk Type, Chunk Flags and Chunk Length
- Verify that association exists
- Sequence of Chunk Types (INIT, INIT ACK, etc)
- Timer checking
- Four way handshake checking
- Heartbeat mechanism
- Protection against INIT/ACK flood DoS attacks, and long-INIT flooding
- Protection against association hijacking

FortiOS also supports SCTP sessions over IPsec VPN tunnels, as well as full traffic and event logging for SCTP sessions.

Protocol port values

The source and destination ports for TCP/UDP/SCTP services are important to get correct. If they are reversed the service will not work. The destination port(s) are the ones that refer to the ports that the computer will be listening on. These are the port numbers that most people are familiar with when they associate a port number to a protocol. In most cases the source port will be one that is randomly assigned by the computer that is not being already used by another service.

Most people associate HTTP with port 80. This means that a web-server will be listening on port 80 for any http requests being sent to the computer. The computer that is sending the request can use any port that is not already assigned to another service or communication session. There are 65,535 ports that it can randomly assign, but because the ports from 1 to 1024 are normally used for listening for incoming communications it is usually not in that range. It is unless there is a specific instance when you know that a communication will be coming from a predefined source port it is best practice to set the source port range from 1 to 65,535.

ICMP

The Internet Control Message Protocol (ICMP) is a protocol layered onto the Internet Protocol Suite to provide error reporting flow control and first-hop gateway redirection. It is normally used by the operating systems of networked computers to send connectivity status query, response and error messages. It is assigned protocol number 1. There is a separate version of the protocol for both IPv4 and for IPv6. It is not designed to be absolutely reliable like TCP.

ICMP is not typically used for transporting data or for end-user network applications with the exception of some diagnostic utilities such as ping and traceroute.

ICMP messages are sent in several situations, for example:

- when a datagram cannot reach its destination,
- time exceeded messages
- redirect messages
- when the gateway does not have the buffering capacity to forward a datagram
- when the gateway can direct the host to send traffic on a shorter route.

Some of the specific ICMP message types are:

- ICMP_ECHO
- ICMP_TIMESTAMP
- ICMP_INFO_REQUEST
- ICMP_ADDRESS

For ICMP error messages, only those reporting an error for an existing session can pass through the firewall. The security policy will allow traffic to be routed, forwarded or denied. If allowed, the ICMP packets will start a new session. Only ICMP error messages of a corresponding security policy is available will be sent back to the source. Otherwise, the packet is dropped. That is, only ICMP packets for a corresponding security policy can traverse the FortiGate unit.

ICMP types and codes

ICMP has a number of messages that are identified by the “Type” field. Some of these types have assigned “Code” fields as well. The table below shows the different types of ICMP Types with their associated codes if there are any.

ICMP types and codes

Type Number	Type Name	Optional Code(s)
0	Echo Reply	
1	Unassigned	
2	Unassigned	

Type Number	Type Name	Optional Code(s)
3	Destination Unreachable	0 Net Unreachable
		1 Host Unreachable
		2 Protocol Unreachable
		3 Port Unreachable
		4 Fragmentation Needed and Don't Fragment was Set
		5 Source Route Failed
		6 Destination Network Unknown
		7 Destination Host Unknown
		8 Source Host Isolated
		9 Communication with Destination Network is Administratively Prohibited
		10 Communication with Destination Host is Administratively Prohibited
		11 Destination Network Unreachable for Type of Service
		12 Destination Host Unreachable for Type of Service
		13 Communication Administratively Prohibited
		14 Host Precedence Violation
		15 Precedence cutoff in effect
4	Source Quench	
5	Redirect	0 Redirect Datagram for the Network (or subnet)
		1 Redirect Datagram for the Host
		2 Redirect Datagram for the Type of Service and Network
		3 Redirect Datagram for the Type of Service and Host
6	Alternate Host Address	
7	Unassigned	

Type Number	Type Name	Optional Code(s)
8	Echo	
9	Router Advertisement	
10	Router Selection	
11	Time Exceeded	0 Time to Live exceeded in Transit 1 Fragment Reassembly Time Exceeded
12	Parameter Problem	0 Pointer indicates the error 1 Missing a Required Option 2 Bad Length
13	Timestamp	
14	Timestamp Reply	
15	Information Request	
16	Information Reply	
17	Address Mask Request	
18	Address Mask Reply	
19	Reserved (for Security)	
20 - 29	Reserved (for Robustness Experiment)	
30	Traceroute	
31	Datagram Conversion Error	
32	Mobile Host Redirect	

Type Number	Type Name	Optional Code(s)
33	IPv6 Where-Are-You	
34	IPv6 I-Am-Here	
35	Mobile Registration	
36	Mobile Registration Reply	
37	Domain Name Request	
38	Domain Name Reply	
39	SKIP	
40	Photuris	
41 - 255	Reserved	

log-invalid-packet

The `log-invalid-packet` CLI setting is one that is intended to log invalid ICMP packets. The exact definition being:

If the FortiGate unit receives an ICMP error packet that contains an embedded `IP(A,B) | TCP(C,D)` header, then if FortiOS can locate the `A:C -> B:D` session it checks to make sure that the sequence number in the TCP header is within the range recorded in the session. If the sequence number is not in range then the ICMP packet is dropped.

When this field is enabled, the FortiGate also log messages that are not ICMP error packets.

Types of logs covered by log-invalid-packet

- Invalid ICMP
 - If ICMP error message verification (see "check-reset-range") is enabled
- Invalid DNS packets
 - DNS packets that contain requests for non-existing domains
- `iprope` check failed
- reverse path check fail
- denied and broadcast traffic
- no session matched

Some other examples of messages that are not errors that will be logged, based on [RFC792](#):

Type 3 messages correspond to "Destination Unreachable Message"

- Type 3, Code 1 = host unreachable
- Type 3, Code 3 = port unreachable

Type 11 messages correspond to "Time Exceeded Message"

- Type 11, Code 0 = time to live exceeded in transit

ICMPv6

Internet Control Message Protocol version 6 (ICMPv6) is the new implementation of the Internet Control Message Protocol (ICMP) that is part of Internet Protocol version 6 (IPv6). The ICMPv6 protocol is defined in [RFC 4443](#).

ICMPv6 is a multipurpose protocol. It performs such things as:

- error reporting in packet processing
- diagnostic functions
- Neighbor Discovery process
- IPv6 multicast membership reporting

It also designed as a framework to use extensions for use with future implementations and changes.

Examples of extensions that have already been written for ICMPv6:

- Neighbor Discovery Protocol (NDP) - a node discovery protocol in IPv6 which replaces and enhances functions of ARP.
- Secure Neighbor Discovery Protocol (SEND) - an extension of NDP with extra security.
- Multicast Router Discovery (MRD) - allows discovery of multicast routers.

ICMPv6 messages use IPv6 packets for transportation and can include IPv6 extension headers. ICMPv6 includes some of the functionality that in IPv4 was distributed among protocols such as ICMPv4, ARP (Address Resolution Protocol), and IGMP (Internet Group Membership Protocol version 3).

ICMPv6 has simplified the communication process by eliminating obsolete messages.

ICMPv6 messages are subdivided into two classes: error messages and information messages.

Error Messages are divided into four categories:

1. Destination Unreachable
2. Time Exceeded
3. Packet Too Big
4. Parameter Problems

Information messages are divided into three groups:

1. Diagnostic messages
2. Neighbor Discovery messages
3. Messages for the management of multicast groups.

ICMPv6 types and codes

ICMPv6 has a number of messages that are identified by the "Type" field. Some of these types have assigned "Code" fields as well. The table below shows the different types of ICMP Types with their associated codes if

there are any.

Type codes 0 – 127 are error messages and type codes 128 – 255 are for information messages.

ICMPv6 types and codes

Type Number	Type Name	Code
0	Reserved	0 - no route to destination
		1 - communication with destination administratively prohibited
		2 - beyond scope of source address
		3 - address unreachable
		4 - port unreachable
		5 - source address failed ingress/egress policy
		6 - reject route to destination
		7 - Error in Source Routing Header
1	Destination Unreachable	
2	Packet Too Big	
3	Time Exceeded	0 - hop limit exceeded in transit
		1 - fragment reassembly time exceeded
4	Parameter Problem	0 - erroneous header field encountered
		1 - unrecognized Next Header type encountered
		2 - unrecognized IPv6 option encountered
100	Private Experimentation	
101	Private Experimentation	
102 - 126	Unassigned	
127	Reserved for expansion if ICMPv6 error messages	

Type Number	Type Name	Code
128	Echo Request	
129	Echo Replay	
130	Multicast Listener Query	
131	Multicast Listener Report	
132	Multicast Listener Done	
133	Router Solicitation	
134	Router Advertisement	
135	Neighbor Solicitation	
136	Neighbor Advertisement	
137	Redirect Message	
138	Router Renumbering	0 - Router Renumbering Command
		1 - Router Renumbering Result
		255 - Sequence Number Reset
139	ICMP Node Information Query	0 - The Data field contains an IPv6 address which is the Subject of this Query.
		1 - The Data field contains a name which is the Subject of this Query, or is empty, as in the case of a NOOP.
		2 - The Data field contains an IPv4 address which is the Subject of this Query.

Type Number	Type Name	Code
140	ICMP Node Information Response	0 - A successful reply. The Reply Data field may or may not be empty.
		1 - The Responder refuses to supply the answer. The Reply Data field will be empty.
		2 - The Qtype of the Query is unknown to the Responder. The Reply Data field will be empty.
141	Inverse Neighbor Discovery Solicitation Message	
142	Inverse Neighbor Discovery Advertisement Message	
143	Version 2 Multicast Listener Report	
144	Home Agent Address Discovery Request Message	
145	Home Agent Address Discovery Reply Message	
146	Mobile Prefix Solicitation	
147	Mobile Prefix Advertisement	
148	Certification Path Solicitation Message	
149	Certification Path Advertisement Message	

Type Number	Type Name	Code
150	ICMP messages utilized by experimental mobility protocols such as Seamoby	
151	Multicast Router Advertisement	
152	Multicast Router Solicitation	
153	Multicast Router Termination	
154	FMIPv6 Messages	
155	RPL Control Message	
156	ILNPv6 Locator Update Message	
157	Duplicate Address Request	
158	Duplicate Address Confirmation	
159 – 199	Unassigned	
200	Private experimentation	
201	Private experimentation	
255	Reserved for expansion of ICMPv6 informational messages	

IP

Internet Protocol (IP) is the primary part of the Network Layer of the OSI Model that is responsible for routing traffic across network boundaries. It is the protocol that is responsible for addressing. IPv4 is probable the version that most people are familiar with and it has been around since 1974. IPv6 is its current successor and due to a

shortage of available IPv4 addresses compared to the explosive increase in the number of devices that use IP addresses, IPv6 is rapidly increasing in use.

When IP is chosen as the protocol type the available option to further specify the protocol is the protocol number. This is used to narrow down which protocol within the Internet Protocol Suite and provide a more granular control.

Protocol number

IP is responsible for more than the address that it is most commonly associated with and there are a number of associated protocols that make up the Network Layer. While there are not 256 of them, the field that identifies them is a numeric value between 0 and 256.

In the Internet Protocol version 4 (IPv4) [RFC791] there is a field called "Protocol" to identify the next level protocol. This is an 8 bit field. In Internet Protocol version 6 (IPv6) [RFC2460], this field is called the "Next Header" field.

Protocol numbers

#	Protocol	Protocol's Full Name
0	HOPOPT	IPv6 Hop-by-Hop Option
1	ICMP	Internet Control Message Protocol
2	IGMP	Internet Group Management
3	GGP	Gateway-to-Gateway
4	IPv4	IPv4 encapsulation Protocol
5	ST	Stream
6	TCP	Transmission Control Protocol
7	CBT	CBT
8	EGP	Exterior Gateway Protocol
9	IGP	Any private interior gateway (used by Cisco for their IGRP)
10	BBN-RCC-MON	BBN RCC Monitoring
11	NVP-II	Network Voice Protocol
12	PUP	PUP
13	ARGUS	ARGUS
14	EMCON	EMCON
15	XNET	Cross Net Debugger

#	Protocol	Protocol's Full Name
16	CHAOS	Chaos
17	UDP	User Datagram Protocol
18	MUX	Multiplexing
19	DCN-MEAS	DCN Measurement Subsystems
20	HMP	Host Monitoring
21	PRM	Packet Radio Measurement
22	XNS-IDP	XEROX NS IDP
23	TRUNK-1	Trunk-1
24	TRUNK-2	Trunk-2
25	LEAF-1	Leaf-1
26	LEAF-2	Leaf-2
27	RDP	Reliable Data Protocol
28	IRTP	Internet Reliable Transaction
29	ISO-TP4	ISO Transport Protocol Class 4
30	NETBLT	Bulk Data Transfer Protocol
31	MFE-NSP	MFE Network Services Protocol
32	MERIT-INP	MERIT Internodal Protocol
33	DCCP	Datagram Congestion Control Protocol
34	3PC	Third Party Connect Protocol
35	IDPR	Inter-Domain Policy Routing Protocol
36	XTP	XTP
37	DDP	Datagram Delivery Protocol
38	IDPR-CMTP	IDPR Control Message Transport Proto
39	TP++	TP++ Transport Protocol

#	Protocol	Protocol's Full Name
40	IL	IL Transport Protocol
41	IPv6	IPv6 encapsulation
42	IPv6	SDRPSource Demand Routing Protocol
43	IPv6-Route	Routing Header for IPv6
44	IPv6-Frag	Fragment Header for IPv6
45	IDRP	Inter-Domain Routing Protocol
46	RSVP	Reservation Protocol
47	GRE	General Routing Encapsulation
48	DSR	Dynamic Source Routing Protocol
49	BNA	BNA
50	ESP	Encap Security Payload
51	AH	Authentication Header
52	I-NLSP	Integrated Net Layer Security TUBA
53	SWIPE	IP with Encryption
54	NARP	NBMA Address Resolution Protocol
55	MOBILE	IP Mobility
56	TLSP	Transport Layer Security Protocol using Kryptonet key management
57	SKIP	SKIP
58	IPv6-ICMP	ICMP for IPv6
59	IPv6-NoNxt	No Next Header for IPv6
60	IPv6-Opts	Destination Options for IPv6
61		any host internal protocol
62	CFTP	CFTP
63		any local network

#	Protocol	Protocol's Full Name
64	SAT-EXPAK	SATNET and Backroom EXPAK
65	KRYPTOLAN	Kryptolan
66	RVD	MIT Remote Virtual Disk Protocol
67	IPPC	Internet Pluribus Packet Core
68		any distributed file system
69	SAT-MON	SATNET Monitoring
70	VISA	VISA Protocol
71	IPCV	Internet Packet Core Utility
72	CPNX	Computer Protocol Network Executive
73	CPHB	Computer Protocol Heart Beat
74	WSN	Wang Span Network
75	PVP	Packet Video Protocol
76	BR-SAT-MON	Backroom SATNET Monitoring
77	SUN-ND	SUN ND PROTOCOL-Temporary
78	WB-MON	WIDEBAND Monitoring
79	WB-EXPAK	WIDEBAND EXPAK
80	ISO-IP	ISO Internet Protocol
81	VMTP	VMTP
82	SECURE-VMTP	SECURE-VMTP
83	VINES	VINES
84	TTP	TTP
84	IPTM	Protocol Internet Protocol Traffic
85	NSFNET-IGP	NSFNET-IGP
86	DGP	Dissimilar Gateway Protocol

#	Protocol	Protocol's Full Name
87	TCF	TCF
88	EIGRP	EIGRP
89	OSPFIGP	OSPFIGP
90	Sprite-RPC	Sprite RPC Protocol
91	LARP	Locus Address Resolution Protocol
92	MTP	Multicast Transport Protocol
93	AX.25	AX.25 Frames
94	IPIP	IP-within-IP Encapsulation Protocol
95	MICP	Mobile Internetworking Control Pro.
96	SCC-SP	Semaphore Communications Sec. Pro.
97	ETHERIP	Ethernet-within-IP Encapsulation
98	ENCAP	Encapsulation Header
99		any private encryption scheme
100	GMTP	GMTP
101	IFMP	Ipsilon Flow Management Protocol
102	PNNI	PNNI over IP
103	PIM	Protocol Independent Multicast
104	ARIS	ARIS
105	SCPS	SCPS
106	QNX	QNX
107	A/N	Active Networks
108	IPComp	IP Payload Compression Protocol
109	SNP	Sitara Networks Protocol
110	Compaq-Peer	Compaq Peer Protocol

#	Protocol	Protocol's Full Name
111	IPX-in-IP	IPX in IP
112	VRRP	Virtual Router Redundancy Protocol
113	PGM	PGM Reliable Transport Protocol
114		any 0-hop protocol
115	L2TP	Layer Two Tunneling Protocol
116	DDX	D-II Data Exchange (DDX)
117	IATP	Interactive Agent Transfer Protocol
118	STP	Schedule Transfer Protocol
119	SRP	SpectraLink Radio Protocol
120	UTI	UTI
121	SMP	Simple Message Protocol
122	SM	SM
123	PTP	Performance Transparency Protocol
124	ISIS over IPv4	
125	FIRE	
126	CRTP	Combat Radio Transport Protocol
127	CRUDP	Combat Radio User Datagram
128	SSCOPMCE	
129	IPLT	
130	SPS	Secure Packet Shield
131	PIPE	Private IP Encapsulation within IP
132	SCTP	Stream Control Transmission Protocol
133	FC	Fibre Channel
134	RSVP-E2E-IGNORE	

#	Protocol	Protocol's Full Name
135	Mobility Header	
136	UDPLite	
137	MPLS-in-IP	
138	manet	
139	HIP	
140	Shim6	
141	WESP	
142	ROHC	
143 – 252	Unassigned	Unassigned
253		Use for experimentation and testing
254		Use for experimentation and testing
255	Reserved	

Further information can be found by researching RFC 5237.

Protocol number

IP is responsible for more than the address that it is most commonly associated with and there are a number of associated protocols that make up the Network Layer. While there are not 256 of them, the field that identifies them is a numeric value between 0 and 256.

In the Internet Protocol version 4 (IPv4) [RFC791] there is a field called “Protocol” to identify the next level protocol. This is an 8 bit field. In Internet Protocol version 6 (IPv6) [RFC2460], this field is called the “Next Header” field.

Protocol numbers

#	Protocol	Protocol's Full Name
0	HOPOPT	IPv6 Hop-by-Hop Option
1	ICMP	Internet Control Message Protocol
2	IGMP	Internet Group Management
3	GGP	Gateway-to-Gateway

#	Protocol	Protocol's Full Name
4	IPv4	IPv4 encapsulation Protocol
5	ST	Stream
6	TCP	Transmission Control Protocol
7	CBT	CBT
8	EGP	Exterior Gateway Protocol
9	IGP	Any private interior gateway (used by Cisco for their IGRP)
10	BBN-RCC-MON	BBN RCC Monitoring
11	NVP-II	Network Voice Protocol
12	PUP	PUP
13	ARGUS	ARGUS
14	EMCON	EMCON
15	XNET	Cross Net Debugger
16	CHAOS	Chaos
17	UDP	User Datagram Protocol
18	MUX	Multiplexing
19	DCN-MEAS	DCN Measurement Subsystems
20	HMP	Host Monitoring
21	PRM	Packet Radio Measurement
22	XNS-IDP	XEROX NS IDP
23	TRUNK-1	Trunk-1
24	TRUNK-2	Trunk-2
25	LEAF-1	Leaf-1
26	LEAF-2	Leaf-2
27	RDP	Reliable Data Protocol

#	Protocol	Protocol's Full Name
28	IRTP	Internet Reliable Transaction
29	ISO-TP4	ISO Transport Protocol Class 4
30	NETBLT	Bulk Data Transfer Protocol
31	MFE-NSP	MFE Network Services Protocol
32	MERIT-INP	MERIT Internodal Protocol
33	DCCP	Datagram Congestion Control Protocol
34	3PC	Third Party Connect Protocol
35	IDPR	Inter-Domain Policy Routing Protocol
36	XTP	XTP
37	DDP	Datagram Delivery Protocol
38	IDPR-CMTP	IDPR Control Message Transport Proto
39	TP++	TP++ Transport Protocol
40	IL	IL Transport Protocol
41	IPv6	IPv6 encapsulation
42	IPv6	SDRPSource Demand Routing Protocol
43	IPv6-Route	Routing Header for IPv6
44	IPv6-Frag	Fragment Header for IPv6
45	IDRP	Inter-Domain Routing Protocol
46	RSVP	Reservation Protocol
47	GRE	General Routing Encapsulation
48	DSR	Dynamic Source Routing Protocol
49	BNA	BNA
50	ESP	Encap Security Payload
51	AH	Authentication Header

#	Protocol	Protocol's Full Name
52	I-NLSP	Integrated Net Layer Security TUBA
53	SWIPE	IP with Encryption
54	NARP	NBMA Address Resolution Protocol
55	MOBILE	IP Mobility
56	TLSP	Transport Layer Security Protocol using Kryptonnet key management
57	SKIP	SKIP
58	IPv6-ICMP	ICMP for IPv6
59	IPv6-NoNxt	No Next Header for IPv6
60	IPv6-Opts	Destination Options for IPv6
61		any host internal protocol
62	CFTP	CFTP
63		any local network
64	SAT-EXPAK	SATNET and Backroom EXPAK
65	KRYPTOLAN	Kryptolan
66	RVD	MIT Remote Virtual Disk Protocol
67	IPPC	Internet Pluribus Packet Core
68		any distributed file system
69	SAT-MON	SATNET Monitoring
70	VISA	VISA Protocol
71	IPCV	Internet Packet Core Utility
72	CPNX	Computer Protocol Network Executive
73	CPHB	Computer Protocol Heart Beat
74	WSN	Wang Span Network
75	PVP	Packet Video Protocol

#	Protocol	Protocol's Full Name
76	BR-SAT-MON	Backroom SATNET Monitoring
77	SUN-ND	SUN ND PROTOCOL-Temporary
78	WB-MON	WIDEBAND Monitoring
79	WB-EXPAK	WIDEBAND EXPAK
80	ISO-IP	ISO Internet Protocol
81	VMTP	VMTP
82	SECURE-VMTP	SECURE-VMTP
83	VINES	VINES
84	TTP	TTP
84	IPTM	Protocol Internet Protocol Traffic
85	NSFNET-IGP	NSFNET-IGP
86	DGP	Dissimilar Gateway Protocol
87	TCF	TCF
88	EIGRP	EIGRP
89	OSPFIGP	OSPFIGP
90	Sprite-RPC	Sprite RPC Protocol
91	LARP	Locus Address Resolution Protocol
92	MTP	Multicast Transport Protocol
93	AX.25	AX.25 Frames
94	IPIP	IP-within-IP Encapsulation Protocol
95	MICP	Mobile Internetworking Control Pro.
96	SCC-SP	Semaphore Communications Sec. Pro.
97	ETHERIP	Ethernet-within-IP Encapsulation
98	ENCAP	Encapsulation Header

#	Protocol	Protocol's Full Name
99		any private encryption scheme
100	GMTP	GMTP
101	IFMP	Ipsilon Flow Management Protocol
102	PNNI	PNNI over IP
103	PIM	Protocol Independent Multicast
104	ARIS	ARIS
105	SCPS	SCPS
106	QNX	QNX
107	A/N	Active Networks
108	IPComp	IP Payload Compression Protocol
109	SNP	Sitara Networks Protocol
110	Compaq-Peer	Compaq Peer Protocol
111	IPX-in-IP	IPX in IP
112	VRRP	Virtual Router Redundancy Protocol
113	PGM	PGM Reliable Transport Protocol
114		any 0-hop protocol
115	L2TP	Layer Two Tunneling Protocol
116	DDX	D-II Data Exchange (DDX)
117	IATP	Interactive Agent Transfer Protocol
118	STP	Schedule Transfer Protocol
119	SRP	SpectraLink Radio Protocol
120	UTI	UTI
121	SMP	Simple Message Protocol
122	SM	SM

#	Protocol	Protocol's Full Name
123	PTP	Performance Transparency Protocol
124	ISIS over IPv4	
125	FIRE	
126	CRTP	Combat Radio Transport Protocol
127	CRUDP	Combat Radio User Datagram
128	SSCOPMCE	
129	IPLT	
130	SPS	Secure Packet Shield
131	PIPE	Private IP Encapsulation within IP
132	SCTP	Stream Control Transmission Protocol
133	FC	Fibre Channel
134	RSVP-E2E-IGNORE	
135	Mobility Header	
136	UDPLite	
137	MPLS-in-IP	
138	manet	
139	HIP	
140	Shim6	
141	WESP	
142	ROHC	
143 – 252	Unassigned	Unassigned
253		Use for experimentation and testing
254		Use for experimentation and testing
255	Reserved	

Further information can be found by researching RFC 5237.

VPN policies

At one point, if you wanted to have secure digital communications between 2 points a private network would be created. This network would only allow the people that were intended to get the communications on it. This is very straightforward if the 2 points are in the same room or even in the same building. It can all be done physically. If you are supposed to be on the secure network

VPNs are an answer to one of today's biggest concerns, how to make digital communications secure between to points that must communicate over the Internet which anybody can have access to.

There are two types of VPNs supported by FortiOS, SSL and IPsec. They are differentiated by the security protocol suites that are used to secure the traffic. These are both described in more detail in the VPN section, but the IPsec VPN can be configured as an **Action** with a firewall policy.

IPsec policies

IPsec policies allow IPsec VPN traffic access to the internal network from a remote location. These policies include authentication information that authenticates users and user group or groups. These policies specify the following:

- the FortiGate firewall interface that provides the physical connection to the remote VPN gateway, usually an interface connected to the Internet
- the FortiGate firewall interface that connects to the private network
- IP addresses associated with data that has to be encrypted and decrypted
- optional: a schedule that restricts when the VPN can operate, and services (or types of data) that can be sent.

For a route-based (interface mode) VPN, you do not configure an IPsec security policy. Instead, you configure two regular ACCEPT security policies, one for each direction of communication, with the IPsec virtual interface as the source or destination interface, as appropriate.

DSRI

The Disable Server Response Inspection (DSRI) options is available for configuration in the CLI. This is used to assist performance when only URL filtering is being used. This allows the system to ignore the HTTP server responses. The setting is configured to be disabled by default.

CLI syntax for changing the status of the DSRI setting

In IPv4 or IPv6 firewall policies

```
config firewall policy|policy6
  edit 0
    set dsri enable|disable
  end
```

In IPv4 or IPv6 interface policies

```
config firewall interface-policy|interface-policy6
  edit 0
    set dsri enable|disable
  end
```


When using the sniffer

```
config firewall sniffer
edit 0
set dsri enable|disable
end
```

Interface policies

Interface policies are implemented before the “security” policies and are only flow based. They are configured in the CLI.

This feature allows you to attach a set of IPS policies with the interface instead of the forwarding path, so packets can be delivered to IPS before entering firewall. This feature is used for following IPS deployments:

- One-Arm: by defining interface policies with IPS and DoS anomaly checks and enabling sniff-mode on the interface, the interface can be used for one-arm IDS;
- IPv6 IPS: IPS inspection can be enabled through interface IPv6 policy. Only IPS signature scan is supported in FortiOS 4.0. IPv6 DoS protection is not supported;
- Scan traffics that destined to FortiGate;
- Scan and log traffics that are silently dropped or flooded by Firewall or Multicast traffic.

IPS sensors can be assigned to an interface policy. Both incoming and outgoing packets are inspected by IPS sensor (signature).

Here is an example of an interface policy,

show full-configuration

```
config firewall interface-policy
edit 1
set status enable
set comments 'test interface policy #1'
set logtraffic utm
set interface "port9"
set srcaddr "all"
set dstaddr "all"
set service "ALL"
set application-list-status disable
set ips-sensor-status disable
set dsri disable
set av-profile-status enable
set av-profile "default"
set webfilter-profile-status disable
set spamfilter-profile-status disable
set dlp-sensor-status disable
set scan-botnet-connections disable
next
end
```

DoS protection

Denial of Service (DoS) policies are primarily used to apply DoS anomaly checks to network traffic based on the FortiGate interface it is entering as well as the source and destination addresses. DoS checks are a traffic anomaly detection feature to identify network traffic that does not fit known or common traffic patterns and

behavior. A common example of anomalous traffic is the denial of service attack. A denial of service occurs when an attacking system starts an abnormally large number of sessions with a target system. The large number of sessions slows down or disables the target system, so that legitimate users can no longer use it.

DoS policies are similar to firewall policies except that instead of defining the way traffic is allowed to flow, they keep track of certain traffic patterns and attributes and will stop traffic displaying those attributes. Further, DoS policies affect only incoming traffic on a single interface. You can further limit a DoS policy by source address, destination address, and service.

DoS configurations have been changed a couple of times in the past. In FortiOS 4.0, DoS protection is moved to the interface policy, so when it is enabled, it is the first thing checked when a packet enters FortiGate. Because of this early detection, DoS policies are a very efficient defense that uses few resources. Denial of service attacks, for example, are detected and its packets dropped before requiring security policy look-ups, antivirus scans, and other protective but resource-intensive operations.

A DoS policy examines network traffic arriving at an interface for anomalous patterns usually indicating an attack. This does not mean that all anomalies experience by the firewall are the result of an intentional attack.

Because an improperly configured DoS anomaly check can interfere with network traffic, no DoS checks are preconfigured on a factory default FortiGate unit. You must create your own before they will take effect. Thresholds for newly created sensors are preset with recommended values that you can adjust to meet the needs of your network.

To create a Denial of Service policy determine if it needs to be an IPv4 or IPv6 policy, then go to:

Policy & Objects > IPv4 DoS Policy for IPv4.

Policy & Objects > IPv6 DoS Policy for IPv6.



The **Enable SSH Deep Scan** feature is enabled by default when creating a new SSL/SSH Inspection profile. There are situations where this feature can cause issues so be sure that you would like it enabled before applying it.

Settings used in configuring DoS

Incoming interface

The interface to which this security policy applies. It will be the that the traffic is coming into the firewall on.

Source address

This will be the address that the traffic is coming from and must be a address listed in the Address section of the Firewall Objects. This can include the predefined “all” address which covers any address coming in on any interface. Multiple addresses or address groups can be chosen

Destination address

This will be the address that the traffic is addressed to. In this case it must be an address that is associated with the firewall itself. For instance it could be one of the interface address of the firewall, a secondary IP address or the interface address assigned to a Virtual IP address. Just like with the Source Address this address must be already configured before being used in the DoS policy. Multiple addresses, virtual IPs or virtual IP groups can be chosen.

Service

While the Service field allows for the use of the ALL service some administrators prefer to optimize the resources of the firewall and only check on the services that will be answered on an interface. Multiple services or service groups can be chosen.

Anomalies

The anomalies can not be configured by the user. They are predefined sensors set up for specific patterns of anomalous traffic

The anomalies that have been predefined for use in the DoS Policies are:

Anomaly Name	Description	Recommended Threshold
tcp_syn_flood	If the SYN packet rate of new TCP connections, including retransmission, to one destination IP address exceeds the configured threshold value, the action is executed.	2000 packets per second.
tcp_port_scan	If the SYN packet rate of new TCP connections, including retransmission, from one source IP address exceeds the configured threshold value, the action is executed.	1000 packets per second.
tcp_src_session	If the number of concurrent TCP connections from one source IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
tcp_dst_session	If the number of concurrent TCP connections to one destination IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
udp_flood	If the UDP traffic to one destination IP address exceeds the configured threshold value, the action is executed.	2000 packets per second.
udp_scan	If the number of UDP sessions originating from one source IP address exceeds the configured threshold value, the action is executed.	2000 packets per second.
udp_src_session	If the number of concurrent UDP connections from one source IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
udp_dst_session	If the number of concurrent UDP connections to one destination IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.

Anomaly Name	Description	Recommended Threshold
icmp_flood	If the number of ICMP packets sent to one destination IP address exceeds the configured threshold value, the action is executed.	250 packets per second.
icmp_sweep	If the number of ICMP packets originating from one source IP address exceeds the configured threshold value, the action is executed.	100 packets per second.
icmp_src_session	If the number of concurrent ICMP connections from one source IP address exceeds the configured threshold value, the action is executed.	300 concurrent sessions
icmp_dst_session	If the number of concurrent ICMP connections to one destination IP address exceeds the configured threshold value, the action is executed.	3000 concurrent sessions
ip_src_session	If the number of concurrent IP connections from one source IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
ip_dst_session	If the number of concurrent IP connections to one destination IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions.
sctp_flood	If the number of SCTP packets sent to one destination IP address exceeds the configured threshold value, the action is executed.	2000 packets per second
sctp_scan	If the number of SCTP sessions originating from one source IP address exceeds the configured threshold value, the action is executed.	1000 packets per second
sctp_src_session	If the number of concurrent SCTP connections from one source IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions
sctp_dst_session	If the number of concurrent SCTP connections to one destination IP address exceeds the configured threshold value, the action is executed.	5000 concurrent sessions

Status

The status field is enabled to enable the sensor for the associated anomaly. In terms of actions performed there is no difference between disabling a sensor and having the action as "Pass" but by disabling sensors that are not being used for blocking or logging you can save some resources of the firewall that can be better used elsewhere.

Logging

Regardless of whether the traffic is blocked or passed through the anomalous traffic will be logged.

Pass

Allows the anomalous traffic to pass through unimpeded.

Block

For Thresholds based on the number of concurrent sessions blocking the anomaly will not allow more than the number of concurrent sessions set as the threshold.

For rate based thresholds where the threshold is measured in packets per second, the Action setting “Block” prevents the overwhelming of the firewall by anomalous traffic in one of 2 ways. Setting which of those 2 ways will be issued is determined in the CLI.

- continuous - blocks packets once an anomaly is detected. This overrides individual anomaly settings.
- periodical - allows matching anomalous traffic up to the rate set by the threshold.



If the period for a particular anomaly is 60 seconds, such as those where the threshold is measured in concurrent sessions, after the 60 second timer has expired, the number of allowed packets that match the anomaly criteria is reset to zero. This means that if you allow 10 sessions through before blocking, after the 60 seconds is up, another 10 will be allowed. The attrition of sessions from expiration should keep the allowed sessions from reaching the maximum.

To set the type of block action for the rate based anomaly sensors:

```
config ips global
    set anomaly-mode continuous
    set anomaly-mode periodical
end
```

Threshold

The threshold can be either in terms of concurrent session or in packets per second depending on which sensor is being referred to.

Quarantine

The quarantine feature is found in the CLI. This setting is used to block any further traffic from a source address that is now considered to be a malicious actor or a source of traffic dangerous to the network. Not only is no more traffic accepted for the duration of the quarantine through the DoS policy but the source IP address of the traffic is added to the banned source ip list. This list is kept in the kernel and used by

- Antivirus
- Data Leak Prevention (DLP)
- Denial of Service (DoS)
- Intrusion Prevention System (IPS)

Any policies that use any of these features will block traffic from the attacker's IP address.

Syntax

```
config firewall {DoS-policy|DoS-policy6}
    edit <policyid>
        set quarantine {none|attacker}
        set quarantine-exipiry {string}
```

```

    set quarantine-log {enable|disable}
end

```

Option	Description
quarantine	Quarantine method. <ul style="list-style-type: none"> <code>none</code> - Quarantine is disabled. <code>attacker</code> - Block all traffic sent from the attacker's IP address. The quarantined IP address is also added to the banned ip list. The destination address is not affected.
quarantine-expiry	Duration of quarantine The format is <code>###d##h##m</code> , ranging from 1 minute to 364 days, 23 hours, and 59 minutes starting from now. The default is <code>0d0h5m</code> . Requires quarantine set to <code>attacker</code> .
quarantine-log	Enables or disables the logging of quarantine events.

One-Arm IDS

Interface-based policy only defines what and how IPS functions are applied to the packets transmitted by the interface. It works no matter if the port is used in a forwarding path or used as an One-Arm device.

To enable One-Arm IDS, the user should first enable sniff-mode on the interface,

```

config system interface
    edit port2
        set ips-sniffer-mode enable
    next
end

```

Once sniff-mode is turned on, both incoming and outgoing packets will be dropped after IPS inspections. The port can be connected to a hub or a switch's SPAN port. Any packet picked up by the interface will still follow the interface policy so different IPS and DoS anomaly checks can be applied.

IPv6 IPS

IPv6 IPS signature scan can be enabled by interface policy. The user can create a normal IPS sensor and assign it to the IPv6 interface policy.

```

config firewall interface-policy6
    edit 1
        set interface "port1"
        set srcaddr6 "all"
        set dstaddr6 "all"
        set service6 "ANY"
        set ips-sensor-status enable
        set ips-sensor "all_default"
    next
end

```

Traffic destined to the FortiGate unit

IPS enabled in firewall policies can only inspect the traffic pass through FortiGate unit, not the traffic destined to FortiGate unit. Enabling IPS in interface-policy allows IPS to pick up any packet on the interface so it is able to inspect attacks targeting FGT.

Dropped, flooded, broadcast, multicast and L2 packets

In many evaluation or certification tests, FortiGate firewall is often required to log any packets dropped by the firewall. In most of cases, these packets are of invalid headers so firewall just drops them silently. It is natural to forward all these packets to IPS first so FortiGate firewall is able to generate logs for invalid packets.

Flooded, broadcast and multicast traffics do not reach any of services in the forwarding path. They can be inspected by the interface policy as long as they match the addresses defined. Potentially, L2 packets can also be sent to IPS for inspection through interface-policy, but it is not enabled in FortiOS 4.0.

GUI and CLI

Now in FortiGate, there are two places that IPS can be enabled, in a firewall policy and in an interface policy. In the firewall policy implementation, IPS sensor can be configured in both CLI and GUI. When adding an IPS sensor to an interface policy it must be done through the CLI. There is no GUI input window for the “Interface Policy”. There is however, a DoS Policy section in the GUI.

Local-In policies

On the FortiGate unit, there are a number of protocols and traffic that is specific to the internal workings of FortiOS. For many of these traffic sources, you can identify a specific port/IP address for this self-originating traffic. The following traffic can be configured to a specific port/IP address:

- SNMP
- Syslog
- alert email
- FortiManager connection IP
- FortiGuard services
- FortiAnalyzer logging
- NTP
- DNS
- Authorization requests such as RADIUS
- FSSO

Security policies control the flow of traffic through the FortiGate unit. The FortiGate unit also includes the option of controlling internal traffic, that is, management traffic.

Each interface includes an allow access configuration to allow management access for specific protocols. Local policies are set up automatically to allow all users all access. Local-in policies takes this a step further, to enable or restrict the user with that access. This also extends beyond the allow access selection.

Local-in policies are configured in the CLI with the commands:

```
config firewall local-in-policy
  edit <policy_number>
    set intf <source_interface>
    set srcaddr <source_address>
```

```

    set dstaddr <destination_address>
    set action {accept | deny}
    set service <service name>
    set schedule <schedule_name>
    set comments <string>
end

```

For example, you can configure a local-in policy so that only administrators can access the FortiGate unit on weekends from a specific management computer at 192.168.21.12, represented by the address object `mgmt-comp1`, using SSH on port 3 (192.168.21.77 represented by the address object `FG-port3`) using the Weekend schedule which defines the time the of access.

```

config firewall local-in-policy
edit <l>
    set intf port3
    set srcaddr mgmt-comp1
    set dstaddr FG-port3
    set action accept
    set service SSH
    set schedule Weekend
end

```

You can also disable a policy should there be a requirement to turn off a policy for troubleshooting or other purpose. To disable a policy enter the commands:

```

config firewall local-in-policy
edit <policy_number>
    set status disable
end

```

Use the same commands with a status of enable to use the policy again.

It is also an option to dedicate the interface as HA management interface by using the setting:

```
set ha-mgmt-intf-only enable
```

Local-in policies are also supported for IPv6 by entering the command:

```
config firewall local-in-policy6.
```



While there is a section under **Policy & Objects** for viewing the existing **Local In Policy** configuration, policies cannot be created or edited here in the GUI. The Local In policies can only be created or edited in the CLI.

Security policy 0

Any security policy that is automatically added by the FortiGate unit has a policy ID number of zero (0). The most common reasons the FortiGate unit creates this policy is:

- The IPsec policy for FortiAnalyzer (and FortiManager version 3.0) is automatically added when an IPsec connection to the FortiAnalyzer unit or FortiManager is enabled.
- The policy to allow FortiGuard servers to be automatically added has a policy ID number of zero.
- The (default) drop rule that is the last rule in the policy and that is automatically added has a policy ID number of zero.
- When a network zone is defined within a VDOM, the intra-zone traffic set to allow or block is managed by policy 0 if it is not processed by a configured security policy.

This policy can appear in logs but will never appear in the security policy list, and therefore, can never be repositioned in the list.

When viewing the FortiGate firewall logs, you may find a log field entry indicating `policyid=0`. The following log message example indicates the log field `policyid=0` in bold.

```
2008-10-06 00:13:49 log_id=0022013001 type=traffic subtype=violation pri=warning
vd=root SN=179089 duration=0 user=N/A group=N/A rule=0 policyid=0 proto=17
service=137/udp app_type=N/A status=deny src=10.181.77.73 srcname=10.181.77.73
dst=10.128.1.161 dstname=10.128.1.161 src_int=N/A dst_int="Internal" sent=0 rcvd=0
src_port=137 dst_port=137 vpn=N/A tran_ip=0.0.0.0 tran_port=0
```

DNS traffic in NGFW policy-mode

FortiOS has an option to enable the creation of an implicit policy to allow DNS traffic.

Certain Application Control profiles may not work properly if DNS traffic is not allowed. Enabling the `implicit-allow-dns` option adds an implicit policy to allow the DNS traffic. This policy is situated in the policy sequence just above the implicit deny policy. Since this is a `config system settings` command, this option can be enabled per VDOM.

CLI

```
config system settings
set implicit-allow-dns {enable|disable}
end
```

Deny policies

Deny security policies deny traffic that is coming into the network. The FortiGate unit automatically blocks traffic that is associated with a deny security policy.

Deny security policies are usually configured when you need to restrict specific traffic, for example, SSH traffic. Deny security policies can also help when you want to block a service, such as DNS, but allow a specific DNS server.

Accept policies

Accept security policies accept traffic that is coming into the network. These policies allow traffic through the FortiGate unit, where the packets are scanned, translated if NAT is enabled, and then sent out to its destination.

Accept security policies are the most common security policies that are created in FortiOS. These security policies are basic policies, such as allowing Internet access, as well as complex policies, such as IPsec VPN.

Fixed port

Some network configurations do not operate correctly if a NAT policy translates the source port of packets used by the connection. NAT translates source ports to keep track of connections for a particular service.

From the CLI you can enable `fixedport` when configuring a security policy for NAT policies to prevent source port translation.

```
config firewall policy
edit <policy-id>
...
```

```
set fixedport enable
...
end
```

However, enabling fixedport means that only one connection can be supported through the firewall for this service. To be able to support multiple connections, add an IP pool, and then select Dynamic IP pool in the policy. The firewall randomly selects an IP address from the IP pool and assigns it to each connection. In this case, the number of connections that the firewall can support is limited by the number of IP addresses in the pool.

Endpoint security

Endpoint security enforces the use of the FortiClient End Point Security (FortiClient and FortiClient Lite) application on your network. It can also allow or deny endpoints access to the network based on the application installed on them.

By applying endpoint security to a security policy, you can enforce this type of security on your network. FortiClient enforcement can check that the endpoint is running the most recent version of the FortiClient application, that the antivirus signatures are up-to-date, and that the firewall is enabled. An endpoint is usually often a single PC with a single IP address being used to access network services through a FortiGate unit.

With endpoint security enabled on a policy, traffic that attempts to pass through, the FortiGate unit runs compliance checks on the originating host on the source interface. Non-compliant endpoints are blocked. If someone is browsing the web, the endpoints are redirected to a web portal which explains the non-compliance and provides a link to download the FortiClient application installer. The web portal is already installed on the FortiGate unit, as a replacement message, which you can modify if required.

Endpoint Security requires that all hosts using the security policy have the FortiClient Endpoint Security agent installed. Currently, FortiClient Endpoint Security is available for Microsoft Windows 2000 and later only.

For more information about endpoint security, see the Security Profiles chapter in the FortiOS Handbook.

Traffic logging

When you enable logging on a security policy, the FortiGate unit records the scanning process activity that occurs, as well as whether the FortiGate unit allowed or denied the traffic according to the rules stated in the security policy. This information can provide insight into whether a security policy is working properly, as well as if there needs to be any modifications to the security policy, such as adding traffic shaping for better traffic performance.

Depending on what the FortiGate unit has in the way of resources, there may be advantages in optimizing the amount of logging taking places. This is why in each policy you are given 3 options for the logging:

- **Disable Log Allowed Traffic** - Does not record any log messages about traffic accepted by this policy.

If you enable Log Allowed Traffic, the following two options are available:

- **Security Events** - This records only log messages relating to security events caused by traffic accepted by this policy.
- **All Sessions** - This records all log messages relating to all of the traffic accepted by this policy.

Depending on the model, if the Log all Sessions option is selected there may be 2 additional options. These options are normally available in the GUI on the higher end models such as the FortiGate 600C or larger.

- **Generate Logs when Session Starts**
- **Capture Packets**

You can also use the CLI to enter the following command to write a log message when a session starts:

```
config firewall policy
  edit <policy-index>
    set logtraffic-start
  end
```

Traffic is logged in the traffic log file and provides detailed information that you may not think you need, but do. For example, the traffic log can have information about an application used (web: HTTP.Image), and whether or not the packet was SNAT or DNAT translated. The following is an example of a traffic log message.

```
2011-04-13
05:23:47
log_id=4
type=traffic
subtype=other
pri=notice
vd=root
status="start"
src="10.41.101.20"
srcname="10.41.101.20"
src_port=58115
dst="172.20.120.100"
dstname="172.20.120.100"
dst_country="N/A"
dst_port=137
tran_ip="N/A"
tran_port=0
tran_sip="10.31.101.41"
tran_sport=58115
service="137/udp"
proto=17
app_type="N/A"
duration=0
rule=1
policyid=1
sent=0
rcvd=0
shaper_drop_sent=0
shaper_drop_rcvd=0
perip_drop=0
src_int="internal"
dst_int="wan1"
SN=97404 app="N/A"
app_cat="N/A"
carrier_ep="N/A"
```

If you want to know more about logging, see the Logging and Reporting chapter in the FortiOS Handbook. If you want to know more about traffic log messages, see the FortiGate Log Message Reference.

IPv6

Internet Protocol version 6 (IPv6) will succeed IPv4 as the standard networking protocol of the Internet. IPv6 provides a number of advances over IPv4 but the primary reason for its replacing IPv4 is its limitation in addresses. IPv4 uses 32 bit addresses which means there is a theoretical limit of 2 to the power of 32. The IPv6 address scheme is based on a 128 bit address or a theoretical limit of 2 to the power of 128.

IPv6 addressing

Possible addresses:

- IPv4 = 4,294,967,296 (over 4 billion)
- IPv6 = 340,282,366,920,938,463,463,374,607,431,768,211,456 (over 340 undecillion - We had to look that term up. We didn't know what a number followed by 36 digits was either)

Assuming a world population of approximately 8 billion people, IPv6 would allow for each individual to have approximately 42,535,295,865,117,200,000,000,000,000 devices with an IP address. That's 42 quintillion devices.

There is little likelihood that you will ever need to worry about these numbers as any kind of serious limitation in addressing but they do give an idea of the scope of the difference in the available addressing.

IPv6 address syntax

Aside from the difference of possible addresses there is also the different formatting of the addresses that will need to be addressed.

A computer would view an IPv4 address as a 32 bit string of binary digits made up of 1s and 0s, broken up into 4 octets of 8 digits separated by a period “.”

Example:

```
10101100.00010000.11111110.00000001
```

To make number more user friendly for humans we translate this into decimal, again 4 octets separated by a period “.” which works out to:

```
172.16.254.1
```

A computer would view an IPv6 address as a 128 bit string of binary digits made up of 1s and 0s, broken up into 8 octets of 16 digits separated by a colon “:”

```
1000000000000001:0000110110111000:101011000001000:1111111000000001:0000000000000000  
0:0000000000000000:0000000000000000:0000000000000000
```

To make number a little more user friendly for humans we translate this into hexadecimal, again 8 octets separated by a colon “:” which works out to:

```
8001:0DB8:AC10:FE01:0000:0000:0000:0000:
```

Because any four-digit group of zeros within an IPv6 address may be reduced to a single zero or altogether omitted, this address can be shortened further to:

8001:0DB8:AC10:FE01:0:0:0:0

or

8001:0DB8:AC10:FE01::

IPv6 packet structure

Each IPv6 packet consists of a mandatory fixed header and optional extension headers, and carries a payload, which is typically either a datagram and/or Transport Layer information. The payload could also contain data for the Internet Layer or Link Layer. Unlike IPv4, IPv6 packets aren't fragmented by routers, requiring hosts to implement Maximum Transmission Unit (MTU) Path Discovery for MTUs larger than the smallest MTU (which is 1280 octets).

Jumbograms and jumbo payloads

In IPv6, packets which exceed the MTU of the underlying network are labeled jumbograms, which consist of a jumbo payload. A jumbogram typically exceeds the IP MTU size limit of 65,535 octets, and provides the jumbo payload option, which can allow up to nearly 4GiB of payload data, as defined in [RFC 2675](#). When the MTU is determined to be too large, the receiving host sends a 'Packet too Big' ICMPv6 type 2 message to the sender.

Fragmentation and reassembly

As noted, packets that are too large for the MTU require hosts to perform MTU Path Discovery to determine the maximum size of packets to send. Packets that are too large require a 'Fragment' extension header, to divide the payload into segments that are 8 octets in length (except for the last fragment, which is smaller). Packets are reassembled according to the extension header and the fragment offset.

Benefits of IPv6

In addition to the expanded number of addresses, some of the other benefits of IPv6 include:

- More efficient routing
- Reduced management requirement
- Stateless auto-reconfiguration of hosts
- Improved methods to change Internet Service Providers
- Better mobility support
- Multi-homing
- Security
- Scoped address: link-local, site-local and global address space

IPv6 in FortiOS

From an administrative point of view IPv6 works almost the same as IPv4 in FortiOS. The primary differences are the use of IPv6 format for addresses and fewer address types for IPv6. There is also no need for NAT if the FortiGate firewall is the interface between IPv6 networks. If the subnets attached to the FortiGate firewall are IPv6 and IPv4 NAT can be configured between the 2 different formats. This will involve either configuring a dual stack routing or IPv4 tunneling configuration. The reason for this is simple. NAT was developed primarily for the purpose of extending the number of usable IPv4 addresses. IPv6's addressing allows for enough available addresses so the NAT is no longer necessary.

When configuring IPv6 in FortiOS, you can create a dual stack route or IPv4-IPv6 tunnel. A dual stack routing configuration implements dual IP layers, supporting both IPv4 and IPv6, in both hosts and routers. An IPv4-IPv6 tunnel is essentially similar, creating a tunnel that encapsulates IPv6 packets within IPv4 headers that carry these IPv6 packets over IPv4 tunnels. The FortiGate unit can also be easily integrated into an IPv6 network. Connecting the FortiGate unit to an IPv6 network is exactly the same as connecting it to an IPv4 network, the only difference is that you are using IPv6 addresses.

By default the IPv6 settings are not displayed in the Web-based Manager. It is just a matter of enabling the display of these feature to use them through the web interface. To enable them just go to **System > Feature Select** and select **IPv6**. Once enabled, you will be able to use IPv6 addresses as well as the IPv4 addressing for the following FortiGate firewall features:

- Static routing
- Policy Routing
- Packet and network sniffing
- Dynamic routing (RIPv6, BGP4+, and OSPFv3)
- IPsec VPN
- DNS
- DHCP
- SSL VPN
- Network interface addressing
- Security Profiles protection
- Routing access lists and prefix lists
- NAT/Route and transparent mode
- NAT 64 and NAT 66
- IPv6 tunnel over IPv4 and IPv4 tunnel over IPv6
- Logging and reporting
- Security policies
- SNMP
- Authentication
- Virtual IPs and groups
- IPv6 over SCTP
- IPv6-specific troubleshooting, such as ping6

IPv6 features

In order to configure IPv6 features using the web-based manager, IPv6 must be enabled using Feature Select. Go to **System > Config > Features**, enable IPv6, and click **Apply**.

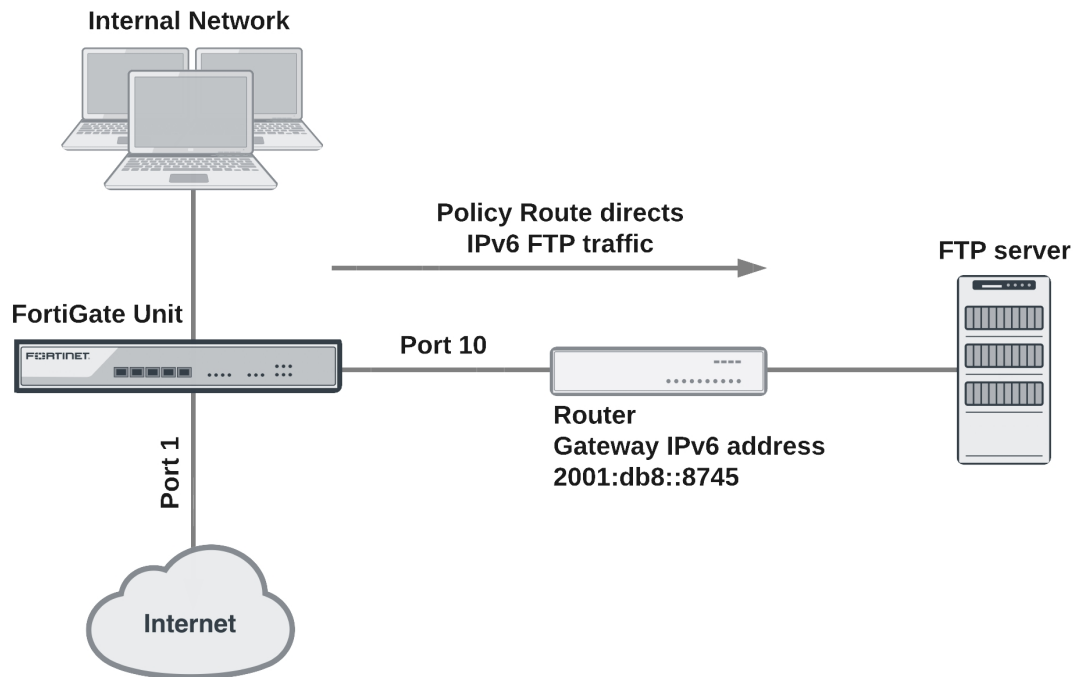
The following IPv6 features are available from the FortiOS web manager:

IPv6 policies

IPv6 security policies are created both for an IPv6 network and a transitional network. A transitional network is a network that is transitioning over to IPv6 but must still have access to the Internet or must connect over an IPv4 network.

These policies allow for this specific type of traffic to travel between the IPv6 and IPv4 networks. The IPv6 options for creating these policies is hidden by default. You must enable this feature under **System > Config > Features**.

IPv6 policy route



IPv6 policy routing

IPv6 policy routing functions in the same way as IPv4 policy routing. To add an IPv6 policy route, go to **Network > Policy Routes** and select **Create New > IPv6 Policy Route**.

Adding an IPv6 Policy route

New Routing Policy

If incoming traffic matches:

Protocol TCP UDP SCTP ANY Specify 0

Incoming interface Ednet (SSID: Student-net)

Source address / mask ::/0

Destination address / mask ::/0

Type of Service Bit Pattern 0x00 Bit Mask 0x00

Then:

Outgoing interface Click to set...

Gateway Address ::

Comments 0/255

OK Cancel

You can also use the following command to add IPv6 policy routes:

```
config router policy6
  edit 0
    set input-device <interface>
    set src <ipv6_ip>
    set dst <ipv6_ip>
    set protocol <0-255>
    set gateway <ipv6_ip>
    set output-device <interface>
    set tos <bit_pattern>
    set tos-mask <bit_mask>
  end
```

IPv6 security policies

IPv6 security policies support all the features supported by IPv4 security policies:

- Policy types and subtypes.
- NAT support including using the destination interface IP address, fixed port, and dynamic IP pools.
- All security features (antivirus, web filtering, application control, IPS, email filtering, DLP, VoIP, and ICAP).
- All traffic shaping options, including: shared traffic shaping, reverse shared traffic shaping, and per-IP traffic shaping.
- All user and device authentication options.

IPv6 explicit web proxy

You can use the explicit web proxy for IPv6 traffic. To do this you need to:

- Enable the IPv6 explicit web proxy from the CLI.
- Enable the explicit web proxy for one or more FortiGate interfaces. These interfaces also need IPv6 addresses.
- Add IPv6 web proxy security policies to allow the explicit web proxy to accept IPv6 traffic.

Use the following steps to set up a FortiGate unit to accept IPv6 traffic for the explicit web proxy at the Internal interface and forward IPv6 explicit proxy traffic out the wan1 interface to the Internet.

1. Enter the following CLI command to enable the IPv6 explicit web proxy:

```
config web-proxy explicit
  set status enable
  set ipv6-status enable
end
```

2. Go to **Network > Interfaces** and edit the **internal** interface, select **Enable Explicit Web Proxy** and select **OK**.
3. Go to **Policy & Objects > Proxy Policy** and select **Create New** to add an IPv6 explicit web proxy security policy with the following settings shown.

This IPv6 explicit web proxy policy allows traffic from all IPv6 IP addresses to connect through the explicit web proxy and through the wan1 interface to any IPv6 addresses that are accessible from the wan1 interface.



If you have enabled both the IPv4 and the IPv6 explicit web proxy, you can combine IPv4 and IPv6 addresses in a single explicit web proxy policy to allow both IPv4 and IPv6 traffic through the proxy.

Example IPv6 Explicit Web Proxy security policy

New Policy	
Explicit Proxy Type	<input checked="" type="radio"/> Web <input type="radio"/> FTP
Enabled On	internal ?
Source Address	<input type="text" value="all"/> +
Outgoing Interface	<input type="text" value="wan1"/> +
Destination Address	<input type="text" value="all"/> +
Schedule	<input type="text" value="always"/>
Action	<input checked="" type="text" value="ACCEPT"/>

Restricting the IP address of the explicit IPv6 web proxy

You can use the following command to restrict access to the IPv6 explicit web proxy using only one IPv6 address. The IPv6 address that you specify must be the IPv6 address of an interface that the explicit HTTP proxy is enabled on. You might want to use this option if the explicit web proxy is enabled on an interface with multiple IPv6 addresses.

For example, to require users to connect to the IPv6 address 2001:db8:0:2::30 to connect to the explicit IPv6 HTTP proxy, use the following command:

```
config web-proxy explicit
  set incoming-ipv6 2001:db8:0:2::30
```

```
end
```

Restricting the outgoing source IP address of the IPv6 explicit web proxy

You can use the following command to restrict the source address of outgoing web proxy packets to a single IPv6 address. The IP address that you specify must be the IPv6 address of an interface that the explicit HTTP proxy is enabled on. You might want to use this option if the explicit HTTP proxy is enabled on an interface with multiple IPv6 addresses.

For example, to restrict the outgoing packet source address to 2001:db8:0:2::50:

```
config http-proxy explicit
  set outgoing-ip6 2001:db8:0:2::50
end
```

VIP64

VIP64 policies can be used to configure static NAT virtual IPv6 address for IPv4 addresses. VIP64 can be configured from the CLI using the following commands:

```
config firewall vip64
  edit <zname_str>
    set arp-reply {enable | disable}
    set color <color_int>
    set comment <comment_str>
    set extip <address_ipv6>[-address_ipv6]
    set extport <port_int>
    set id <id_num_str>
    set mappedip [<start_ipv4>-<end_ipv4>]
    set mappedport <port_int>
    set portforward {enable | disable}
    set src-filter <addr_str>
  end
```

VIP64 CLI Variables and Defaults

Variable	Description	Default
<zname_str>	Enter the name of this virtual IP address.	No default.
arp-reply {enable disable}	Select to respond to ARP requests for this virtual IP address.	enable
color <color_int>	Enter the number of the color to use for the group icon in the web-based manager.	0
comment <comment_str>	Enter comments relevant to the configured virtual IP.	No default.

Variable	Description	Default
<code>extip <address_ipv6>[-address_ipv6]</code>	<p>Enter the IP address or address range on the external interface that you want to map to an address or address range on the destination network.</p> <p>If <code>mappedip</code> is an IP address range, the FortiGate unit uses <code>extip</code> as the first IP address in the external IP address range, and calculates the last IP address required to create an equal number of external and mapped IP addresses for one-to-one mapping.</p> <p>To configure a dynamic virtual IP that accepts connections destined for any IP address, set <code>extip</code> to <code>::</code>.</p>	<code>::</code>
<code>extport <port_int></code>	<p>Enter the external port number that you want to map to a port number on the destination network.</p> <p>This option only appears if <code>portforward</code> is enabled.</p> <p>If <code>portforward</code> is enabled and you want to configure a static NAT virtual IP that maps a range of external port numbers to a range of destination port numbers, set <code>extport</code> to the first port number in the range. Then set <code>mappedport</code> to the start and end of the destination port range. The FortiGate unit automatically calculates the end of the <code>extport</code> port number range.</p>	0
<code>id <id_num_str></code>	Enter a unique identification number for the configured virtual IP. Not checked for uniqueness. Range 0 - 65535.	No default.

Variable	Description	Default
mappedip [<start_ipv4>-<end_ipv4>]	<p>Enter the IP address or IP address range on the destination network to which the external IP address is mapped.</p> <p>If <code>mappedip</code> is an IP address range, the FortiGate unit uses <code>extip</code> as the first IP address in the external IP address range, and calculates the last IP address required to create an equal number of external and mapped IP addresses for one-to-one mapping.</p> <p>If <code>mappedip</code> is an IP address range, the FortiGate unit uses <code>extip</code> as a single IP address to create a one-to-many mapping.</p>	0.0.0.0
mappedport <port_int>	<p>Enter the port number on the destination network to which the external port number is mapped.</p> <p>You can also enter a port number range to forward packets to multiple ports on the destination network.</p> <p>For a static NAT virtual IP, if you add a map to port range the FortiGate unit calculates the external port number range.</p>	0
portforward {enable disable}	Select to enable port forwarding. You must also specify the port forwarding mappings by configuring <code>extport</code> and <code>mappedport</code> .	disable
src-filter <addr_str>	Enter a source address filter. Each address must be in the form of an IPv4 subnet (x:x:x:x:x:x/n). Separate addresses with spaces.	null

VIP46 policies can be used to configure static NAT virtual IPv4 address for IPv6 addresses. VIP46 can be configured from the CLI using the following commands (see the table below for variable details):

```
config firewall vip46
  edit <name_str>
    set arp-reply {enable | disable}
    set color <color_int>
    set comment <comment_str>
    set extip <address_ipv4>[-address_ipv4]
    set extport <port_int>
```

```

    set id <id_num_str>
    set mappedip [<start_ipv6>--<end_ipv6>]
    set mappedport <port_int>
    set portforward {enable | disable}
    set src-filter <add_str>
end

```

VIP46 CLI Variables and Defaults

Variable	Description	Default
<name_str>	Enter the name of this virtual IP address.	No default.
arp-reply {enable disable}	Select to respond to ARP requests for this virtual IP address.	enable
color <color_int>	Enter the number of the color to use for the group icon in the web-based manager.	0
comment <comment_str>	Enter comments relevant to the configured virtual IP.	No default.
extip <address_ipv4>[- address_ipv4]	<p>Enter the IP address or address range on the external interface that you want to map to an address or address range on the destination network.</p> <p>If <code>mappedip</code> is an IP address range, the FortiGate unit uses <code>extip</code> as the first IP address in the external IP address range, and calculates the last IP address required to create an equal number of external and mapped IP addresses for one-to-one mapping.</p> <p>To configure a dynamic virtual IP that accepts connections destined for any IP address, set <code>extip</code> to 0.0.0.0.</p>	0.0.0.0

Variable	Description	Default
<code>extport <port_int></code>	<p>Enter the external port number that you want to map to a port number on the destination network.</p> <p>This option only appears if <code>portforward</code> is enabled.</p> <p>If <code>portforward</code> is enabled and you want to configure a static NAT virtual IP that maps a range of external port numbers to a range of destination port numbers, set <code>extport</code> to the first port number in the range. Then set <code>mappedport</code> to the start and end of the destination port range. The FortiGate unit automatically calculates the end of the <code>extport</code> port number range.</p>	0
<code>id <id_num_str></code>	Enter a unique identification number for the configured virtual IP. Not checked for uniqueness. Range 0 - 65535.	No default.
<code>mappedip [<start_ipv6>-<end_ipv6>]</code>	<p>Enter the IP address or IP address range on the destination network to which the external IP address is mapped.</p> <p>If <code>mappedip</code> is an IP address range, the FortiGate unit uses <code>extip</code> as the first IP address in the external IP address range, and calculates the last IP address required to create an equal number of external and mapped IP addresses for one-to-one mapping.</p> <p>If <code>mappedip</code> is an IP address range, the FortiGate unit uses <code>extip</code> as a single IP address to create a one-to-many mapping.</p>	::

Variable	Description	Default
<code>mappedport <port_int></code>	<p>Enter the port number on the destination network to which the external port number is mapped.</p> <p>You can also enter a port number range to forward packets to multiple ports on the destination network.</p> <p>For a static NAT virtual IP, if you add a map to port range the FortiGate unit calculates the external port number range.</p>	0
<code>portforward {enable disable}</code>	Select to enable port forwarding. You must also specify the port forwarding mappings by configuring <code>extport</code> and <code>mappedport</code> .	disable
<code>src-filter <addr_str></code>	Enter a source address filter. Each address must be in the form of an IPv4 subnet (x.x.x.x/n). Separate addresses with spaces.	null

IPv6 network address translation

NAT66, NAT64, and DNS64 are now supported for IPv6. These options provide IPv6 NAT and DNS capabilities with IPv6-IPv4 tunneling or dual stack configurations. The commands are available only in the CLI.

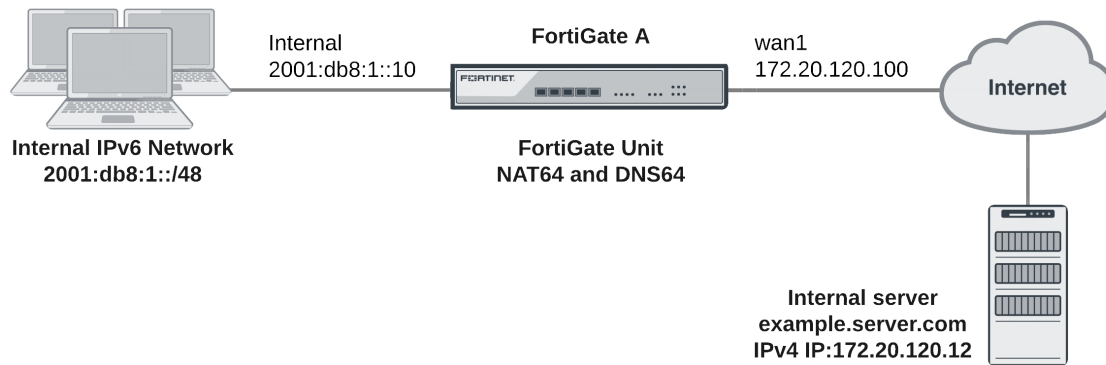
Fortinet supports all features described in [RFC 6146](#). However, for DNS64 there is no support for handling Domain Name System Security Extensions (DNSSEC). DNSSEC is for securing types of information that are provided by the DNS as used on an IP network or networks. You can find more information about DNS64 in [RFC 6147](#).

NAT64 and DNS64 (DNS proxy)

NAT64 is used to translate IPv6 addresses to IPv4 addresses so that a client on an IPv6 network can communicate transparently with a server on an IPv4 network.

NAT64 is usually implemented in combination with the DNS proxy called DNS64. DNS64 synthesizes AAAA records from A records and is used to synthesize IPv6 addresses for hosts that only have IPv4 addresses. 'DNS proxy' and 'DNS64' are interchangeable terms.

Example NAT64 configuration



With a NAT64 and DNS64 configuration in place on a FortiGate unit, clients on an IPv6 network can transparently connect to addresses on an IPv4 network. NAT64 and DNS64 perform the IPv4 to IPv6 transition, allowing clients that have already switched to IPv6 addresses to continue communicating with servers that still use IPv4 addresses.

To enable NAT64 and DNS64, use the following CLI commands:

Enable NAT64

```
config system nat64
  set status enable
end
```

Enable the DNS proxy on the IPv6 interface

```
config system dns-server
  edit internal
  end
```

In your DHCP6 configuration, configure the IPv6 interface IP address as the DNS6 server IP address. The FortiGate will proxy DNS requests to the system DNS server.

```
config system dhcp6 server
  edit 1
    set interface internal
    config ip-range
      edit 1
        set start-ip 2001:db8:1::11
        set end-ip 2001:db8:1::20
      end
    set dns-server1 2001:db8:1::10
  end
```

NAT64 policies

You can configure security policies for NAT64 using the web-based manager. For these options to appear, the feature must be enabled using **System > Feature Visibility**. You can then configure the policies under **Policy & Objects > NAT64 Policy**.

NAT64 policies and can also be configured from the CLI using the following command:

```
config firewall policy64
```

In the following section, you will configure a NAT64 policy that allows connections from an internal IPv6 network to an external IPv4 network.

Configuring NAT64 to allow a host on the IPv6 network to connect to the Internet server

In this example, the Internal IPv6 network address is 2001:db8:1::/48 and the external IPv4 network address is 172.20.120.0/24. NAT64 is configured to allow a user on the internal network to connect to the server at IPv4 address 172.20.120.12. In this configuration, sessions exiting the wan1 interface must have their source address changed to an IPv4 address in the range 172.20.120.200 to 172.20.120.210.

Enter the following command to enable NAT64:

```
config system nat64
  set status enable
end
```

Enabling NAT64 with the `config system nat64` command means that all IPv6 traffic received by the current VDOM can be subject to NAT64 if the source and destination address matches an NAT64 security policy.

By default, the setting `always-synthesize-aaaa-record` is enabled. If you disable this setting, the DNS proxy (DNS64) will attempt to find an AAAA records for queries to domain names and therefore resolve the host names to IPv6 addresses. If the DNS proxy cannot find an AAAA record, it synthesizes one by adding the NAT64 prefix to the A record.

By using the `nat64-prefix` option of the `config system nat64` command to change the default nat64 prefix from the well-known prefix of 64:ff9b::/96 and setting `always-synthesize-aaaa-record` to enable (default), the DNS proxy does not check for AAAA records but rather synthesizes AAAA records.

As an alternative to the above entry, there is the optional configuration that would allow the resolution of CNAME queries.

```
config system nat64
  set status enable
  set nat64-prefix 64:ff9b::/96
  set always-synthesize-aaaa-record enable
end
```

Enter the following command to add an IPv6 firewall address for the internal network:

```
config firewall address6
  edit internal-net6
    set ip6 2001:db8:1::/48
  end
```

Enter the following command to add an IPv4 firewall address for the external network:

```
config firewall address
  edit external-net4
    set subnet 172.20.120.0/24
    set associated-interface wan1
  end
```

Enter the following command to add an IP pool containing the IPv4 address that the should become the source address of the packets exiting the wan1 interface:

```
config firewall ippool
```

```
edit exit-pool4
    set startip 172.20.120.200
    set endip 172.20.120.210
end
```

Enter the following command to add a NAT64 policy that allows connections from the internal IPv6 network to the external IPv4 network:

```
config firewall policy64
    edit 0
        set srcintf internal
        set srcaddr internal-net6
        set dstintf wan1
        set dstaddr external-net4
        set action accept
        set schedule always
        set service ANY
        set logtraffic enable
        set ippool enable
        set poolname exit-pool4
    end
```

The `srcaddr` can be any IPv6 firewall address and the `dstaddr` can be any IPv4 firewall address.

Other NAT64 policy options include `fixedport`, which can be used to prevent NAT64 from changing the destination port. You can also configure traffic shaping for NAT64 policies.

How a host on the internal IPv6 network communicates with `example.server.com` that only has IPv4 address on the Internet

1. The host on the internal network does a DNS lookup for `example.server.com` by sending a DNS query for an AAAA record for `example.server.com`.
2. The DNS query is intercepted by the FortiGate DNS proxy.
3. The DNS proxy attempts to resolve the query with a DNS server on the Internet and discovers that there are no AAAA records for `example.server.com`.
4. The previous step is skipped if `always-synthesize-aaaa-record` is enabled.
5. The DNS proxy performs an A-record query for `example.server.com` and gets back an RRSet containing a single A record with the IPv4 address `172.20.120.12`.
6. The DNS proxy then synthesizes an AAAA record. The IPv6 address in the AAAA record begins with the configured NAT64 prefix in the upper 96 bits and the received IPv4 address in the lower 32 bits. By default, the resulting IPv6 address is `64:ff9b::172.20.120.12`.
7. The host on the internal network receives the synthetic AAAA record and sends a packet to the destination address `64:ff9b::172.20.120.12`.
8. The packet is routed to the FortiGate internal interface where it is accepted by the NAT64 security policy.
9. The FortiGate unit translates the destination address of the packets from IPv6 address `64:ff9b::172.20.120.12` to IPv4 address `172.20.120.12` and translates the source address of the packets to `172.20.120.200` (or another address in the IP pool range) and forwards the packets out the `wan1` interface to the Internet.

NAT66

NAT66 is used for translating an IPv6 source or destination address to a different IPv6 source or destination address. NAT66 is not as common or as important as IPv4 NAT, as many IPv6 addresses do not need NAT66 as

much as IPv4 NAT. However, NAT66 can be useful for a number of reasons. For example, you may have changed the IP addresses of some devices on your network but want traffic to still appear to be coming from their old addresses. You can use NAT66 to translate the source addresses of packets from the devices to their old source addresses.

In FortiOS, NAT66 options can be added to an IPv6 security policy from the CLI. Configuring NAT66 is very similar to configuring NAT in an IPv4 security policy. For example, use the following command to add an IPv6 security policy that translates the source address of IPv6 packets to the address of the destination interface (similar to IPv4 source NAT):

```
config firewall policy6
  edit 0
    set srcintf internal
    set dstintf wan1
    set srcaddr internal_net
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
    set nat enable
  end
```

It's also useful to translate one IPv6 source address to another address that is not the same as the address of the exiting interface. You can do this using IP pools. For example, enter the following command to add an IPv6 IP pool containing one IPv6 IP address:

```
config firewall ippool6
  edit example_6_pool
    set startip 2001:db8::
    set endip 2001:db8::
  end
```

Enter the following command to add an IPv6 firewall address that contains a single IPv6 IP address.

```
config firewall address6
  edit device_address
    set ip6 2001:db8::132/128
  end
```

Enter the following command to add an IPv6 security policy that accepts packets from a device with IP address 2001:db8::132 and translates the source address to 2001:db8::.

```
config firewall policy6
  edit 0
    set srcintf internal
    set dstintf wan1
    set srcaddr device_address
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
    set nat enable
    set ippool enable
    set poolname example_6_pool
  end
```

NAT66 destination address translation

NAT66 can also be used to translate destination addresses. This is done in an IPv6 policy by using IPv6 virtual IPs. For example, enter the following command to add an IPv6 virtual IP that maps the destination address 2001:db8::dd to 2001:db8::ee.

```
config firewall vip6
  edit example-vip6
    set extip 2001:db8::dd
    set mappedip 2001:db8::ee
  end
```

Enter the following command to add an IPv6 security policy that accepts packets with a destination address 2001:db8::dd and translates that destination address to 2001:db8::ee.

```
config firewall policy6
  edit 0
    set srcintf internal
    set dstintf wan1
    set srcaddr all
    set dstaddr example-vip6
    set action accept
    set schedule always
    set service ANY
  end
```

NAT64 and NAT66 session failover

The FortiGate Clustering Protocol (FGCP) supports IPv6, NAT64, and NAT66 session failover. If session pickup is enabled, these sessions are synchronized between cluster members and, after an HA failover, the sessions will resume with only minimal interruption.

NAT46

NAT46 is used to translate IPv4 addresses to IPv6 addresses so that a client on an IPv4 network can communicate transparently with a server on an IPv6 network.

To enable NAT46, use the following CLI command:

```
config firewall vip46
```

NAT46 policies

Security policies for NAT46 can be configured from the web-based manager. For these options to appear in the web-based manager, this feature must be enabled using **System > Feature Visibility**. You can then configure the policies under **Policy & Objects > NAT46 Policy**.

NAT46 policies can also be configured from the CLI using the following command:

```
config firewall policy46
```

IPv6 tunneling

IPv6 Tunneling is the act of tunneling IPv6 packets from an IPv6 network through an IPv4 network to another IPv6 network. This is different than Network Address Translation (NAT) because once the packet reaches its final destination the true originating address of the sender will still be readable. The IPv6 packets are encapsulated

within packets with IPv4 headers, which carry their IPv6 payload through the IPv4 network. This type of configuration is more appropriate for those who have completely transitional over to IPv6, but need an Internet connection, which is still mostly IPv4 addresses.

The key to IPv6 tunneling is the ability of the 2 devices, whether they are a host or a network device, to be dual stack compatible. They have to be able to work with both IPv4 and IPv6 at the same time. In the process the entry node of the tunnel portion of the path will create an encapsulating IPv4 header and transmit the encapsulated packet. The exit node at the end of the tunnel receives the encapsulated packet. The IPv4 header is removed. The IPv6 header is updated and the IPv6 packet is processed.

There are two types of tunnels in IPv6:

Automatic tunnels	Automatic tunnels are configured by using IPv4 address information embedded in an IPv6 address – the IPv6 address of the destination host includes information about which IPv4 address the packet should be tunneled to.
Configured tunnels	Configured tunnels must be configured manually. These tunnels are used when using IPv6 addresses that do not have any embedded IPv4 information. The IPv6 and IPv4 addresses of the endpoints of the tunnel must be specified.

Tunnel configurations

There are a few ways in which the tunneling can be performed depending on which segment of the path between the end points of the session the encapsulation takes place.

Network Device to Network Device	Dual stack capable devices connected by an IPv4 infrastructure can tunnel IPv6 packets between themselves. In this case, the tunnel spans one segment of the path taken by the IPv6 packets.
Host to Network Device	Dual stack capable hosts can tunnel IPv6 packets to an intermediary IPv6 or IPv4 network device that is reachable through an IPv4 infrastructure. This type of tunnel spans the first segment of the path taken by the IPv6 packets.
Host to Host	Dual stack capable hosts that are interconnected by an IPv4 infrastructure can tunnel IPv6 packets between themselves. In this case, the tunnel spans the entire path taken by the IPv6 packets.
Network Device to Host	Dual stack capable network devices can tunnel IPv6 packets to their final destination IPv6 or IPv4 host. This tunnel spans only the last segment of the path taken by the IPv6 packets.

Regardless of whether the tunnel starts at a host or a network device, the node that does the encapsulation needs to maintain soft state information, such as the maximum transmission unit (MTU), about each tunnel in order to process the IPv6 packets.

Use the following command to tunnel IPv6 traffic over an IPv4 network. The IPv6 interface is configured under `config system interface`. The command to do the reverse is `config system ipv6-tunnel`. These commands are not available in transparent mode.

```
config system sit-tunnel
  edit <tunnel name>
    set destination <tunnel _address>
    set interface <name>
```

```

    set ip6 <address_ipv6>
    set source <address_ipv4>
end

```

Variable	Description	Default
edit <tunnel_name>	Enter a name for the IPv6 tunnel.	No default.
destination <tunnel_address>	The destination IPv4 address for this tunnel.	0.0.0.0
interface <name>	The interface used to send and receive traffic for this tunnel.	No default.
ip6 <address_ipv6>	The IPv6 address for this tunnel.	No default.
source <address_ipv4>	The source IPv4 address for this tunnel.	0.0.0.0

Tunneling IPv6 through IPsec VPN

A variation on the tunneling IPv6 through IPv4 is using an IPsec VPN tunnel between two FortiGate devices. FortiOS supports IPv6 over IPsec. In this sort of scenario, 2 networks using IPv6 behind FortiGate units are separated by the Internet, which uses IPv4. An IPsec VPN tunnel is created between the 2 FortiGate units and a tunnel is created over the IPv4 based Internet but the traffic in the tunnel is IPv6. This has the additional advantage of making the traffic secure as well.

For configuration information, see [IPv6 IPsec VPN on page 1](#).

IPv6 support for GRE tunnels

You can use IPv6 addresses can be used at both ends of a GRE tunnel in the same way as with IPv4.

The configuration is similar to how you set up the tunnel for IPv4. However, when you configure the specific tunnel, you need to set the `ip-version` option to 6. This will enable IPv6-specific options for the tunnel.

CLI

```

config system gre-tunnel
  edit <name of tunnel>
    set ip-version 6
    set remote-gw6 <IPv6 address of the remote gateway>
    set local-gw-6 <IPv6 address of the local gateway>
  end


```

SIP over IPv6

FortiOS supports Sessions Initiate Protocol (SIP) over IPv6. The SIP application-level gateway (ALG) can process SIP messages that use IPv6 addresses in the headers, bodies, and in the transport stack. The SIP ALG cannot modify the IPv6 addresses in the SIP headers so FortiGate units cannot perform SIP or RTP NAT over IPv6 and also cannot translate between IPv6 and IPv4 addresses.

In the scenario shown below, a SIP phone connects to the Internet through a FortiGate unit operating. The phone and the SIP and RTP servers all have IPv6 addresses.

The FortiGate unit has IPv6 security policies that accept SIP sessions. The SIP ALG understands IPv6 addresses and can forward IPv6 sessions to their destinations. Using SIP application control features the SIP ALG can also apply rate limiting and other settings to SIP sessions.



To enable SIP support for IPv6 add an IPv6 security policy that accepts SIP packets and includes a VoIP profile.

New Fortinet FortiGate IPv6 MIB fields

The following IPv6 MIB fields have been added to the Fortinet FortiGate MIB. These MIB entries can be used to display IPv6 session and policy statistics.

- IPv6 Session Counters:

```

fgSysSes6Count
fgSysSes6Rate1
fgSysSes6Rate10
fgSysSes6Rate30
fgSysSes6Rate60
  
```

- IPv6 Policy Statistics:

```

fgFwPol6StatsTable
fgFwPol6StatsEntry
FgFwPol6StatsEntry
fgFwPol6ID
fgFwPol6PktCount
fgFwPol6ByteCount
  
```

- IPv6 Session Statistics:

```

fgIp6SessStatsTable
fgIp6SessStatsEntry
FgIp6SessStatsEntry
fgIp6SessNumber
  
```

The `fgSysSesCount` and `fgSysSesRateX` MIBs report statistics for IPv4 plus IPv6 sessions combined. This behavior was not changed.

New OIDs

The following OIDs have been added:

```
FORTINET-FORTIGATE-MIB:fortinet.fnFortiGateMib.fgSystem.fgSystemInfo
.fgSysSes6Count      1.3.6.1.4.1.12356.101.4.1.15
.fgSysSesRate1       1.3.6.1.4.1.12356.101.4.1.16
.fgSysSesRate10      1.3.6.1.4.1.12356.101.4.1.17
.fgSysSesRate30      1.3.6.1.4.1.12356.101.4.1.18
.fgSysSesRate60      1.3.6.1.4.1.12356.101.4.1.19

FORTINET-FORTIGATE-MIB:fortinet.fnFortiGateMib.fgFirewall.fgFwPolicies.fgFwPolTables
.fgFwPol6StatsTable.fgFwPol6StatsEntry.fgFwPol6ID      1.3.6.1.4.1.12356.101.5.1.2.2.1.1
.fgFwPol6StatsTable.fgFwPol6StatsEntry.fgFwPol6PktCount 1.3.6.1.4.1.12356.101.5.1.2.2.1.
2
.fgFwPol6StatsTable.fgFwPol6StatsEntry.fgFwPol6ByteCount 1.3.6.1.4.1.12356.101.5.1.2.2.1
.3

FORTINET-FORTIGATE-MIB:fortinet.fnFortiGateMib.fgInetProto.fgInetProtoTables
.fgIp6SessStatsTable.fgIp6SessStatsEntry.fgIp6SessNumber 1.3.6.1.4.1.12356.101.11.2.3.1.1
```

EXAMPLE SNMP get/walk output

```
// Session6 stats excerpt from sysinfo:
snmpwalk -v2c -cpublic 192.168.1.111 1.3.6.1.4.1.12356.101.4
FORTINET-FORTIGATE-MIB::fgSysSes6Count.0 = Gauge32: 203
FORTINET-FORTIGATE-MIB::fgSysSes6Rate1.0 = Gauge32: 10 Sessions Per Second
FORTINET-FORTIGATE-MIB::fgSysSes6Rate10.0 = Gauge32: 2 Sessions Per Second
FORTINET-FORTIGATE-MIB::fgSysSes6Rate30.0 = Gauge32: 1 Sessions Per Second
FORTINET-FORTIGATE-MIB::fgSysSes6Rate60.0 = Gauge32: 0 Sessions Per Second

// FwPolicy6 table:
snmpwalk -v2c -cpublic 192.168.1.111 1.3.6.1.4.1.12356.101.5.1.2.2
FORTINET-FORTIGATE-MIB::fgFwPol6ID.1.3 = INTEGER: 3
FORTINET-FORTIGATE-MIB::fgFwPol6ID.1.4 = INTEGER: 4
FORTINET-FORTIGATE-MIB::fgFwPol6PktCount.1.3 = Counter64: 4329
FORTINET-FORTIGATE-MIB::fgFwPol6PktCount.1.4 = Counter64: 0
FORTINET-FORTIGATE-MIB::fgFwPol6ByteCount.1.3 = Counter64: 317776
FORTINET-FORTIGATE-MIB::fgFwPol6ByteCount.1.4 = Counter64: 0

// IP6SessNumber:
snmpwalk -v2c -cpublic 192.168.1.111 1.3.6.1.4.1.12356.101.11.2.3.1
FORTINET-FORTIGATE-MIB::fgIp6SessNumber.1 = Counter32: 89
```

IPv6 per-IP traffic shaper

You can add any Per-IP traffic shaper to an IPv6 security policy using the following command:

```
config firewall policy6
edit 0
set per-ip-shaper "new-perip-shaper"
end
```

DHCPv6

You can use DHCP with IPv6 using the CLI. To configure DHCP, ensure IPv6 is enabled by going to **System** > **Feature Visibility** and enabling **IPv6**.

Use the CLI command

```
config system dhcp6
```

For more information on the configuration options, see the [FortiGate CLI Reference](#).

DHCP delegated mode

Downstream IPv6 interfaces can receive address assignments on delegated subnets from a DHCP server that serves an upstream interface.

DHCPv6-PD configuration

Enable DHCPv6 Prefix Delegation on upstream interface (port10):

```
config system interface
  edit "port10"
    config ipv6
      set dhcp6-prefix-delegation enable
    end
  end
```

Assign delegated prefix on downstream interface (port1). Optionally, specific delegated prefixes can be specified:

```
config system interface
  edit "port1"
    config ipv6
      set ip6-mode delegated
      set ip6-upstream-interface "port10"
      set ip6-subnet ::1:0:0:0:1/64
      set ip6-send-adv enable
      config ipv6-delegated-prefix-list
        edit 1
          set upstream-interface "port10"
          set autonomous-flag enable
          set onlink-flag enable
          set subnet 0:0:0:100::/64
        end
      end
    end
  end
```

DHCPv6 server configuration

Configuring a server that uses delegated prefix and DNS from upstream:

```
config system dhcp6 server
  edit 1
    set dns-service delegated
    set interface "wan2"
    set upstream-interface "wan1"
    set ip-mode delegated
    set subnet 0:0:0:102::/64
  end
```

DHCPv6 relay

You can use the following command to configure a FortiGate interface to relay DHCPv6 queries and responses from one network to a network with a DHCPv6 server and back. The command enables DHCPv6 relay and includes adding the IPv6 address of the DHCP server that the FortiGate unit relays DHCPv6 requests to:

```
config system interface
  edit internal
    config ipv6
      set dhcp6-relay-service enable
      set dhcp6-relay-type regular
      set dhcp6-relay-ip 2001:db8:0:2::30
    end
```

IPv6 forwarding

Policies, IPS, Application Control, flow-based antivirus, web filtering, and DLP

FortiOS fully supports flow-based inspection of IPv6 traffic. This includes full support for IPS, application control, virus scanning, and web filtering.

To add flow-based inspection to IPv6 traffic go to **Policy & Objects > IPv6 Policy** and select **Create New** to add an IPv6 Security Policy. Configure the policy to accept the traffic to be scanned. Under **Security Profiles**, select the profiles to apply to the traffic.

Obtaining IPv6 addresses from an IPv6 DHCP server

From the CLI, you can configure any FortiGate interface to get an IPv6 address from an IPv6 DHCP server. For example, to configure the wan2 interface to get an IPv6 address from an IPv6 DHCP server enter the following command:

```
config system interface
  edit wan2
    config ipv6
      set ip6-mode dhcp
    end
```

Authentication support

RADIUS

FortiOS's supports IPv6 RADIUS authentication. When configuring the FortiGate interface and the RADIUS server (under `config system interface` and `config user radius` respectively), the server IP address can be set as IPv6.

Captive portal

Captive portal supports IPv6. It works with remote RADIUS authentication and WiFi interfaces.

FSSO

FortiGate FSSO supports connecting to an FSSO agent over IPv6 and collecting and sending IPv6 details about endpoints. This is all enforced the same way as IPv4 FSSO traffic.

CLI

```
config user fsso
  edit <fsso agent name>
    set source-ip6 <IPv6 address for source>
  end
```



The source-ip6 option accepts either an IPv6 format address or a Fully Qualified Domain Name (FQDN). The FQDN can parse to an IPv6 address.

IPv6 Neighbor Discovery Proxy

This feature provides support for proxying the IPv6 Neighbor Discovery (ND) protocol to allow the forwarding of the following ICMP messages between upstream and downstream interfaces:

- Router Advertisement (RA)
- Neighbor Solicitation (NS)
- Neighbor Advertisement (NA)
- Router Solicitation (RS)
- Redirect



Normally, only one interface will receive RA traffic. This will automatically be considered the upstream interface.

The Neighbor Discovery (ND) protocol is used to discover the Link Layer address of IPv6 destinations. In IPv4, this is achieved by using ARP.

Configure ND Proxy in the CLI using the following syntax:

```
config system nd-proxy
  set status {enable|disable}
  set member <interface> <interface> [<interface>...]
end
```

Option	Description
status	Enable/disable the use of neighbor discovery proxy
member	List of interfaces using the neighbor discovery proxy

An example of a configuration can be found in the IPv6 Configuration section under [IPv6 Neighbor Discovery Proxy on page 663](#)



More information on Neighbor Discovery Proxies (ND Proxy) is available at [RFC 4389](#).

IPv6 configuration

This section contains configuration information for IPv6 on FortiOS. Attempts are made to include scenarios in each section to better assist with the configuration and to orient the information toward a particular task.

You will find information on the following:



By default IPv6 configurations do not appear in the web-based manager. You need to enable the feature first.

To enable IPv6:

1. Go to **System > Features**.
2. Select **IPv6** and click **Apply**.

IPv6 address groups

To create IPv6 address groups from existing IPv6 addresses - web-based manager

Your company has 3 internal servers with IPv6 addresses that it would like to group together for the purposes of a number of policies.

1. Go to **Policy & Objects > Addresses** and select **Create New > Address Group**.
2. Select **IPv6 Group**, and fill out the fields with the following information:

Group Name	Web_Server_Cluster
Members	Web_Server-1
	Web_Server-2
	Web_Server-3

3. Select **OK**.

To create IPv6 address groups from existing IPv6 addresses - CLI

```
config firewall addrgrp6
  edit Web_Server_Cluster
    set member Web_Server-1 Web_Server-2 Web_Server-3
  end
```

To verify that the addresses were added correctly

1. Go to **Policy & Objects > Addresses**. Check that the addresses have been added to the address list and that they are correct.
2. From the CLI, enter the following commands:

```
config firewall addrgrp6
  edit <the name of the address that you wish to verify>
    Show full-configuration
```

IPv6 address ranges

You can configure IPv6 address ranges in both the GUI and the CLI.

To configure IPv6 address ranges - web-based manager:

1. Go to **Policy & Objects > Addresses**.
2. Set the **Type** to **IP Range** and enter the IPv6 addresses as shown:

To configure IPv6 address ranges - CLI:

```
config firewall address6
  edit ipv6range
    set type iprange
    set start-ip 2001:db8:0:2::30
    set end-ip 2001:db8:0:2::31
  end
```

IPv6 firewall addresses

Scenario: Mail server

You need to create an IPv6 address for the Mail Server on Port1 of your internal network. This server is on the network off of port1.

- The IP address is 2001:db8:0:2::20/128
- There should be a tag for this address being for a server.

Configuring the Example using the GUI

1. Go to **Policy & Objects > Objects > Addresses** and select **Create New > Address**.
2. Select **IPv6 Address** and fill out the fields with the following information

Name	Mail_Server
Type	Subnet

IPv6 Address	2001:db8:0:2::20/128
---------------------	----------------------

3. Select **OK**.

Configuring the Example using the CLI

Enter the following CLI command:

```
config firewall address6
edit Mail_Server
set type ipprefix
set subnet 2001:db8:0:2::20/128
end
```

Scenario: First floor network

You need to create an IPv6 address for the subnet of the internal network off of Port1. These computers connect to port1. The network uses the IPv6 addresses: fdde:5a7d:f40b:2e9d:xxxx:xxxx:xxxx:xxxx

There should be a reference to this being the network for the 1st floor of the building.

1. Go to **Policy & Objects > Objects > Addresses**
2. Select **Create New > Address**. Select **IPv6 Address** and fill out the fields with the following information:

Name	Internal_Subnet_1
Type	Subnet / IP Range
IPv6 Address	2001:db8:0:2::/64
Comments	Network for 1st Floor

3. Select **OK**.

4. Enter the following CLI command:

```
config firewall address6
edit Internal_Subnet_1
set comment "Network for 1st Floor"
set type ipprefix
set subnet 2001:db8:0:2::/64
end
```

Scenario: Accounting team

You need to create an IPv6 address for the Accounting Team that's on the 1st Floor. These users are off of various ports of the FortiGate, but they have all been assigned addresses between 2001:db8:0:2::2000 and 2001:db8:0:2::a000

Configuring the example using the GUI

1. Go to **Policy & Objects > Objects > Addresses** and select **Create New > Address**.
2. Select **IPv6 Address** and fill out the fields with the following information

Name	Accounting_Team
Type	IP Range
Subnet / IP Range	2001:db8:0:2::2000-2001:db8:0:2::a000

3. Select **OK**.

Configuring the Example using the CLI

Enter the following CLI command:

```
config firewall address6
edit Accounting_Team
set type iprange
set visibility enable
set start-ip 2001:db8:0:2::2000
set end-ip 2001:db8:0:2::a000
end
```

To verify that the addresses were added correctly:

1. Go to **Policy & Objects > Objects > Addresses**. Check that the addresses have been added to the address list and that they are correct.
2. Enter the following CLI command:

```
config firewall address6
edit <the name of the address that you wish to verify>
Show full-configuration
```

IPv6 SSH

FortiGate supports SSH traffic through IPv6. When the proxy option is set to `ssh` in a proxy policy, IPv6 source and destination address options become available and SSH profiles can be assigned to IPv6 firewall policies.

Syntax in IPv6 firewall policy

```
config firewall policy6
edit 1
set utm-status enable
set ssh-filter-profile <example>
end
```

Syntax in proxy policy

```
config firewall proxy-policy
edit 1
set proxy ssh
set srcaddr6 "all"
set dstaddr6 "all"
end
```

Logging

When a proxy policy is being used, SSH traffic logs are generated by `wad` instead of the kernel.

ICMPv6

The IT Manager is doing some diagnostics and would like to temporarily block the successful replies of ICMP Node information Responses between 2 IPv6 networks.

The ICMP type for ICMP Node information responses is 140. The codes for a successful response is 0.

To configure ICMPv6 - web-based manager:

1. Go to **Policy & Objects > Services** and select **Create New > Service**.
2. Fill out the fields with the following information

Name	diagnostic-test1
Service Type	Firewall
Show in Service List	Enabled
Category	Uncategorized
Protocol Type	ICMP6
Type	140

3. Select **OK**.
4. Enter the following CLI command:

```
config firewall service custom
edit diagnostic-test1
set protocol ICMP6
set icmp-type 140
set icmp-code 0
set visibility enable
end
```

To verify that the category was added correctly:

1. Go to **Policy & Objects > Services**. Check that the services have been added to the services list and that they are correct.
2. Enter the following CLI command:

```
config firewall service custom
edit <the name of the service that you wish to verify>
show full-configuration
```

IPv6 IPsec VPN

This chapter describes how to configure your FortiGate unit's IPv6 IPsec VPN functionality.



By default IPv6 configurations do not appear in the web-based manager. You need to enable the feature first.

To enable IPv6:

1. Go to **System > Features**.
2. Select **IPv6** and click **Apply**.

The topics in this section include:

- [Overview of IPv6 IPsec support](#)
- [Configuring IPv6 IPsec VPNs](#)
- [Site-to-site IPv6 over IPv6 VPN example](#)
- [Site-to-site IPv4 over IPv6 VPN example](#)
- [Site-to-site IPv6 over IPv4 VPN example](#)

Overview of IPv6 IPsec support

FortiOS supports route-based IPv6 IPsec, but not policy-based. This section describes how IPv6 IPsec support differs from IPv4 IPsec support.

Where both the gateways and the protected networks use IPv6 addresses, sometimes called IPv6 over IPv6, you can create either an auto-keyed or manually-keyed VPN. You can also combine IPv6 and IPv4 addressing in an auto-keyed VPN in the following ways:

IPv4 over IPv6	<p>The VPN gateways have IPv6 addresses.</p> <p>The protected networks have IPv4 addresses. The phase 2 configurations at either end use IPv4 selectors.</p>
IPv6 over IPv4	<p>The VPN gateways have IPv4 addresses.</p> <p>The protected networks use IPv6 addresses. The phase 2 configurations at either end use IPv6 selectors.</p>

Compared with IPv4 IPsec VPN functionality, there are some limitations:

- Except for IPv6 over IPv4, remote gateways with Dynamic DNS are not supported.
- Selectors cannot be firewall address names. Only IP address, address range and subnet are supported.
- Redundant IPv6 tunnels are not supported.

Certificates

On a VPN with IPv6 phase 1 configuration, you can authenticate using VPN certificates in which the common name (cn) is an IPv6 address. The `cn-type` keyword of the `user peer` command has an option, `ipv6`, to support this.

Configuring IPv6 IPsec VPNs

Configuration of an IPv6 IPsec VPN follows the same sequence as for an IPv4 route-based VPN: phase 1 settings, phase 2 settings, security policies, and routing.

Phase 1 configuration

In the web-based manager, you define the Phase 1 as IPv6 in the Advanced settings. Enable the IPv6 Version check box. You can then enter an IPv6 address for the remote gateway.

In the CLI, you define an IPsec phase 1 configuration as IPv6 by setting `ip-version` to 6. Its default value is 4. Then, the `local-gw` and `remote-gw` keywords are hidden and the corresponding `local-gw6` and `remote-gw6` keywords are available. The values for `local-gw6` and `remote-gw6` must be IPv6 addresses. For example:

```
config vpn ipsec phase1-interface
edit tunnel6
set ip-version 6
set remote-gw6 0:123:4567::1234
set interface port3
set proposal 3des-md5
end
```

Phase 2 configuration

To create an IPv6 IPsec phase 2 configuration in the web-based manager, you need to define IPv6 selectors in the Advanced settings. Change the default “0.0.0.0/0” address for Source address and Destination address to the IPv6 value “::/0”. If needed, enter specific IPv6 addresses, address ranges or subnet addresses in these fields.

In the CLI, set `src-addr-type` and `dst-addr-type` to `ip6`, `range6` or `subnet6` to specify IPv6 selectors. By default, zero selectors are entered, “::/0” for the `subnet6` address type, for example. The simplest IPv6 phase 2 configuration looks like this:

```
config vpn ipsec phase2-interface
edit tunnel6_p2
set phase1name tunnel6
set proposal 3des-md5
set src-addr-type subnet6
set dst-addr-type subnet6
end
```

Security policies

To complete the VPN configuration, you need a security policy in each direction to permit traffic between the protected network’s port and the IPsec interface. You need IPv6 policies unless the VPN is IPv4 over IPv6.

Routing

Appropriate routing is needed for both the IPsec packets and the encapsulated traffic within them. You need a route, which could be the default route, to the remote VPN gateway via the appropriate interface. You also need a route to the remote protected network via the IPsec interface.

To create a static route - web-based manager:

1. Go to **Network > Static Routes**.

2. Select the drop-down arrow on the **Create New** button and select **IPv6 Route**.
3. Enter the information and select **OK**.

To create a static route - CLI:

1. In the CLI, use the `router static6` command. For example, where the remote network is `fec0:0000:0000:0004::/64` and the IPsec interface is `toB`:

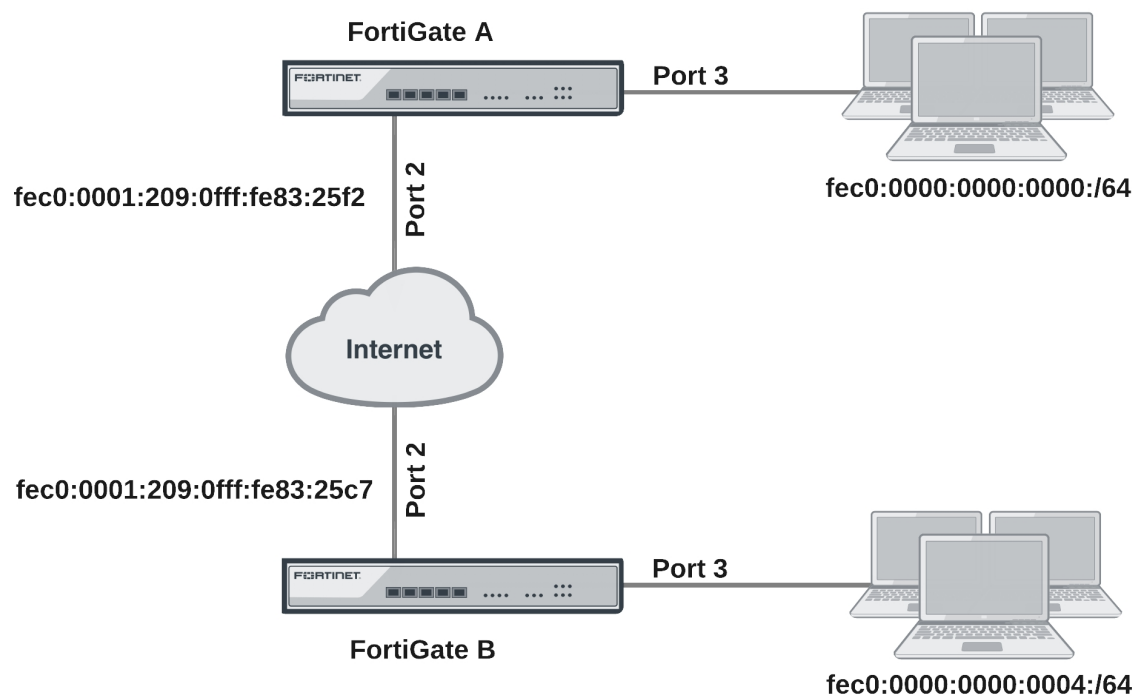
```
config router static6
  edit 1
    set device port2
    set dst 0::/0
  next
  edit 2
    set device toB
    set dst fec0:0000:0000:0004::/64
  next
end
```

If the VPN is IPv4 over IPv6, the route to the remote protected network is an IPv4 route. If the VPN is IPv6 over IPv4, the route to the remote VPN gateway is an IPv4 route.

Site-to-site IPv6 over IPv6 VPN example

In this example, computers on IPv6-addressed private networks communicate securely over public IPv6 infrastructure.

Example IPv6-over-IPv6 VPN topology



Configure FortiGate A interfaces

Port 2 connects to the public network and port 3 connects to the local network.

```
config system interface
  edit port2
    config ipv6
      set ip6-address fec0::0001:209:0fff:fe83:25f2/64
    end
  next
  edit port3
    config ipv6
      set ip6-address fec0::0000:209:0fff:fe83:25f3/64
    end
  next
end
```

Configure FortiGate A IPsec settings

The phase 1 configuration creates a virtual IPsec interface on port 2 and sets the remote gateway to the public IP address FortiGate B. This configuration is the same as for an IPv4 route-based VPN, except that `ip-version` is set to 6 and the `remote-gw6` keyword is used to specify an IPv6 remote gateway address.

```
config vpn ipsec phase1-interface
  edit toB
    set ip-version 6
    set interface port2
    set remote-gw6 fec0:0000:0000:0003:209:0fff:fe83:25c7
    set dpd enable
    set psksecret maryhadalittlelamb
    set proposal 3des-md5 3des-sha1
  end
```

By default, phase 2 selectors are set to accept all subnet addresses for source and destination. The default setting for `src-addr-type` and `dst-addr-type` is `subnet`. The IPv6 equivalent is `subnet6`. The default subnet addresses are 0.0.0.0/0 for IPv4, `::/0` for IPv6.

```
config vpn ipsec phase2-interface
  edit toB2
    set phase1name toB
    set proposal 3des-md5 3des-sha1
    set pfs enable
    set replay enable
    set src-addr-type subnet6
    set dst-addr-type subnet6
  end
```

Configure FortiGate A security policies

Security policies are required to allow traffic between port3 and the IPsec interface toB in each direction. The address `all6` must be defined using the `firewall address6` command as `::/0`.

```
config firewall policy6
  edit 1
    set srcintf port3
    set dstintf toB
    set srcaddr all6
    set dstaddr all6
    set action accept
```

```

        set service ANY
        set schedule always
    next
    edit 2
        set srcintf toB
        set dstintf port3
        set srcaddr all6
        set dstaddr all6
        set action accept
        set service ANY
        set schedule always
    end

```

Configure FortiGate A routing

This simple example requires just two static routes. Traffic to the protected network behind FortiGate B is routed via the virtual IPsec interface toB. A default route sends all IPv6 traffic out on port2.

```

config router static6
    edit 1
        set device port2
        set dst 0::/0
    next
    edit 2
        set device toB
        set dst fec0:0000:0000:0004::/64
    end

```

Configure FortiGate B

The configuration of FortiGate B is very similar to that of FortiGate A. A virtual IPsec interface toA is configured on port2 and its remote gateway is the public IP address of FortiGate A. Security policies enable traffic to pass between the private network and the IPsec interface. Routing ensures traffic for the private network behind FortiGate A goes through the VPN and that all IPv6 packets are routed to the public network.

```

config system interface
    edit port2
        config ipv6
            set ip6-address fec0::0003:209:0fff:fe83:25c7/64
        end
    next
    edit port3
        config ipv6
            set ip6-address fec0::0004:209:0fff:fe83:2569/64
        end
    end
config vpn ipsec phase1-interface
    edit toA
        set ip-version 6
        set interface port2
        set remote-gw6 fec0:0000:0000:0001:209:0fff:fe83:25f2
        set dpd enable
        set psksecret maryhadalittlelamb
        set proposal 3des-md5 3des-sha1
    end
config vpn ipsec phase2-interface
    edit toA2

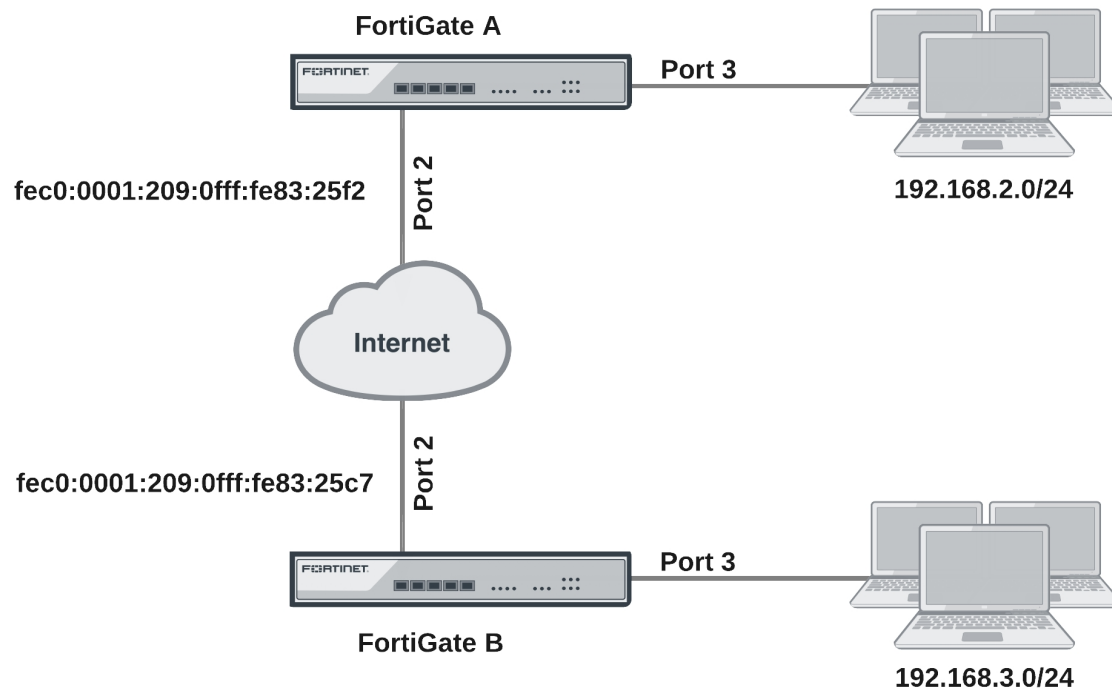
```

```
        set phase1name toA
        set proposal 3des-md5 3des-sha1
        set pfs enable
        set replay enable
        set src-addr-type subnet6
        set dst-addr-type subnet6
    end
config firewall policy6
    edit 1
        set srcintf port3
        set dstintf toA
        set srcaddr all6
        set dstaddr all6
        set action accept
        set service ANY
        set schedule always
    next
    edit 2
        set srcintf toA
        set dstintf port3
        set srcaddr all6
        set dstaddr all6
        set action accept
        set service ANY
        set schedule always
    end
config router static6
    edit 1
        set device port2
        set dst 0::/0
    next
    edit 2
        set device toA
        set dst fec0:0000:0000:0000::/64
end
```

Site-to-site IPv4 over IPv6 VPN example

In this example, two private networks with IPv4 addressing communicate securely over IPv6 infrastructure.

Example IPv4-over-IPv6 VPN topology



Configure FortiGate A interfaces

Port 2 connects to the IPv6 public network and port 3 connects to the IPv4 LAN.

```
config system interface
  edit port2
    config ipv6
      set ip6-address fec0::0001:209:0fff:fe83:25f2/64
    end
  next
  edit port3
    set 192.168.2.1/24
  end
```

Configure FortiGate A IPsec settings

The phase 1 configuration is the same as in the IPv6 over IPv6 example.

```
config vpn ipsec phase1-interface
  edit toB
    set ip-version 6
    set interface port2
    set remote-gw6 fec0:0000:0000:0003:209:0fff:fe83:25c7
    set dpd enable
    set psksecret maryhadalittlelamb
    set proposal 3des-md5 3des-sha1
  end
```

The phase 2 configuration is the same as you would use for an IPv4 VPN. By default, phase 2 selectors are set to accept all subnet addresses for source and destination.

```
config vpn ipsec phase2-interface
edit toB2
    set phaselname toB
    set proposal 3des-md5 3des-sha1
    set pfs enable
    set replay enable
end
```

Configure FortiGate A security policies

Security policies are required to allow traffic between port3 and the IPsec interface toB in each direction. These are IPv4 security policies.

```
config firewall policy
edit 1
    set srcintf port3
    set dstintf toB
    set srcaddr all
    set dstaddr all
    set action accept
    set service ANY
    set schedule always
next
edit 2
    set srcintf toB
    set dstintf port3
    set srcaddr all
    set dstaddr all
    set action accept
    set service ANY
    set schedule always
end
```

Configure FortiGate A routing

This simple example requires just two static routes. Traffic to the protected network behind FortiGate B is routed via the virtual IPsec interface toB using an IPv4 static route. A default route sends all IPv6 traffic, including the IPv6 IPsec packets, out on port2.

```
config router static6
edit 1
    set device port2
    set dst 0::/0
next
edit 2
    set device toB
    set dst 192.168.3.0/24
end
```

Configure FortiGate B

The configuration of FortiGate B is very similar to that of FortiGate A. A virtual IPsec interface toA is configured on port2 and its remote gateway is the public IP address of FortiGate A. The IPsec phase 2 configuration has

IPv4 selectors.

IPv4 security policies enable traffic to pass between the private network and the IPsec interface. An IPv4 static route ensures traffic for the private network behind FortiGate A goes through the VPN and an IPv6 static route ensures that all IPv6 packets are routed to the public network.

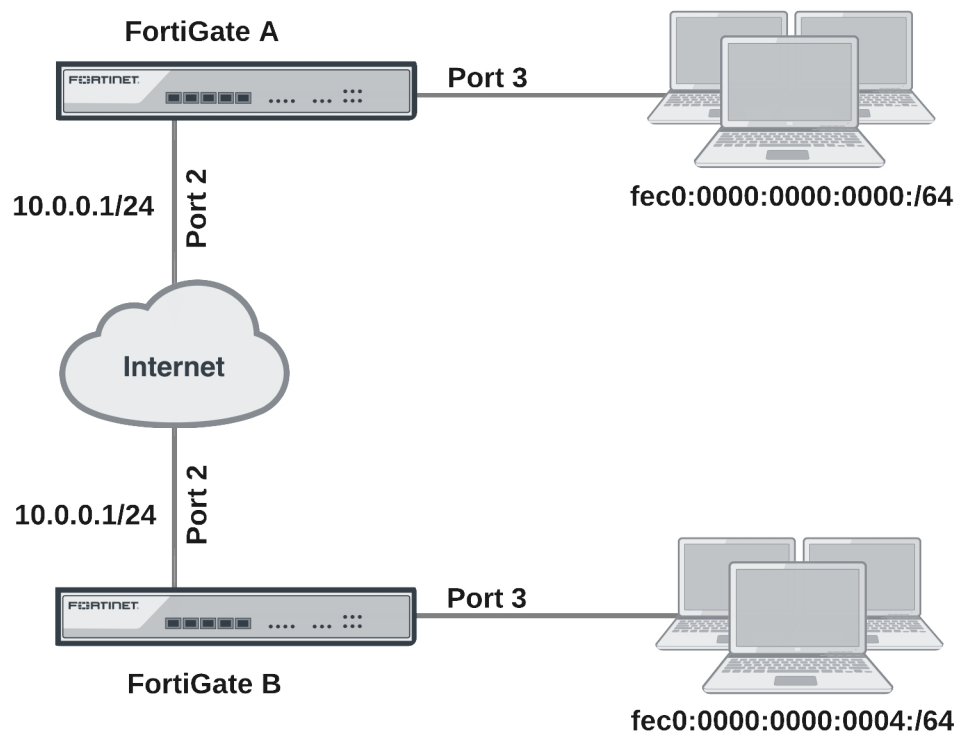
```
config system interface
  edit port2
    config ipv6
      set ip6-address fec0::0003:fe83:25c7/64
    end
  next
  edit port3
    set 192.168.3.1/24
  end
config vpn ipsec phase1-interface
  edit toA
    set ip-version 6
    set interface port2
    set remote-gw6 fec0:0000:0000:0001:209:0fff:fe83:25f2
    set dpd enable
    set psksecret maryhadalittlelamb
    set proposal 3des-md5 3des-sha1
  end
config vpn ipsec phase2-interface
  edit toA2
    set phase1name toA
    set proposal 3des-md5 3des-sha1
    set pfs enable
    set replay enable
  end
config firewall policy
  edit 1
    set srcintf port3
    set dstintf toA
    set srcaddr all
    set dstaddr all
    set action accept
    set service ANY
    set schedule always
  next
  edit 2
    set srcintf toA
    set dstintf port3
    set srcaddr all
    set dstaddr all
    set action accept
    set service ANY
    set schedule always
  end
config router static6
  edit 1
    set device port2
    set dst 0::/0
  next
  edit 2
    set device toA
    set dst 192.168.2.0/24
```

end

Site-to-site IPv6 over IPv4 VPN example

In this example, IPv6-addressed private networks communicate securely over IPv4 public infrastructure.

Example IPv6-over-IPv4 VPN topology



Configure FortiGate A interfaces

Port 2 connects to the IPv4 public network and port 3 connects to the IPv6 LAN.

```
config system interface
  edit port2
    set 10.0.0.1/24
  next
  edit port3
    config ipv6
      set ip6-address fec0::0001:209:0fff:fe83:25f3/64
    end
  end
```

Configure FortiGate A IPsec settings

The phase 1 configuration uses IPv4 addressing.

```
config vpn ipsec phase1-interface
  edit toB
```

```

        set interface port2
        set remote-gw 10.0.1.1
        set dpd enable
        set psksecret maryhadalittlelamb
        set proposal 3des-md5 3des-sha1
    end

```

The phase 2 configuration uses IPv6 selectors. By default, phase 2 selectors are set to accept all subnet addresses for source and destination. The default setting for `src-addr-type` and `dst-addr-type` is `subnet`. The IPv6 equivalent is `subnet6`. The default subnet addresses are `0.0.0.0/0` for IPv4, `::/0` for IPv6.

```

config vpn ipsec phase2-interface
    edit toB2
        set phase1name toB
        set proposal 3des-md5 3des-sha1
        set pfs enable
        set replay enable
        set src-addr-type subnet6
        set dst-addr-type subnet6
    end

```

Configure FortiGate A security policies

IPv6 security policies are required to allow traffic between `port3` and the IPsec interface `toB` in each direction. Define the address `all6` using the `firewall address6` command as `::/0`.

```

config firewall policy6
    edit 1
        set srcintf port3
        set dstintf toB
        set srcaddr all6
        set dstaddr all6
        set action accept
        set service ANY
        set schedule always
    next
    edit 2
        set srcintf toB
        set dstintf port3
        set srcaddr all6
        set dstaddr all6
        set action accept
        set service ANY
        set schedule always
    end

```

Configure FortiGate A routing

This simple example requires just two static routes. Traffic to the protected network behind FortiGate B is routed via the virtual IPsec interface `toB` using an IPv6 static route. A default route sends all IPv4 traffic, including the IPv4 IPsec packets, out on `port2`.

```

config router static6
    edit 1
        set device toB
        set dst fec0:0000:0000:0004::/64
    end
config router static
    edit 1

```

```
set device port2
set dst 0.0.0.0/0
set gateway 10.0.0.254
end
```

Configure FortiGate B

The configuration of FortiGate B is very similar to that of FortiGate A. A virtual IPsec interface toA is configured on port2 and its remote gateway is the IPv4 public IP address of FortiGate A. The IPsec phase 2 configuration has IPv6 selectors.

IPv6 security policies enable traffic to pass between the private network and the IPsec interface. An IPv6 static route ensures traffic for the private network behind FortiGate A goes through the VPN and an IPv4 static route ensures that all IPv4 packets are routed to the public network.

```
config system interface
  edit port2
    set 10.0.1.1/24
  next
  edit port3
    config ipv6
      set ip6-address fec0::0004:209:0fff:fe83:2569/64
    end
config vpn ipsec phase1-interface
  edit toA
    set interface port2
    set remote-gw 10.0.0.1
    set dpd enable
    set psksecret maryhadalittlelamb
    set proposal 3des-md5 3des-sha1
  end
config vpn ipsec phase2-interface
  edit toA2
    set phase1name toA
    set proposal 3des-md5 3des-sha1
    set pfs enable
    set replay enable
    set src-addr-type subnet6
    set dst-addr-type subnet6
  end
config firewall policy6
  edit 1
    set srcintf port3
    set dstintf toA
    set srcaddr all6
    set dstaddr all6
    set action accept
    set service ANY
    set schedule always
  next
  edit 2
    set srcintf toA
    set dstintf port3
    set srcaddr all6
    set dstaddr all6
    set action accept
    set service ANY
    set schedule always
  end
config router static6
  edit 1
    set device toA
```

```

        set dst fec0:0000:0000:0000::/64
    end
config router static
    edit 1
        set device port2
        set gateway 10.0.1.254
    end

```

TCP MSS values

TCP MSS values, which control the maximum amount of data that can be sent in a single packet, can be set for IPv6 policies (for both the sender and the receiver). You can configure TCP MSS values in IPv6 using the following CLI commands:

```

config firewall policy6
    edit <index_int>
        set tcp-mss-sender <value>
        set tcp-mss-receiver <value>
    end

```

BGP and IPv6

FortiGate units support IPv6 over BGP using the same `config router bgp` command as IPv4, but different subcommands.

The main CLI keywords have IPv6 equivalents that are identified by the “6” on the end of the keyword, such as with `config network6` or `set allowas-in6`.

IPv6 BGP commands include:

```

config router bgp
    set activate6 {enable | disable}
    set allowas-in6 <max_num_AS_integer>
    set allowas-in-enable6 {enable | disable}
    set as-override6 {enable | disable}
    set attribute-unchanged6 [as-path] [med] [next-hop]
    set capability-default-originate6 {enable | disable}
    set capability-graceful-restart6 {enable | disable}
    set default-originate-route-map6 <routemap_str>
    set distribute-list-in6 <access-list-name_str>
    set distribute-list-out6 <access-list-name_str>
    set filter-list-in6 <aspath-list-name_str>
    set filter-list-out6 <aspath-list-name_str>
    set maximum-prefix6 <prefix_integer>
    set maximum-prefix-threshold6 <percentage_integer>
    set maximum-prefix-warning-only6 {enable | disable}
    set next-hop-self6 {enable | disable}
    set prefix-list-in6 <prefix-list-name_str>
    set prefix-list-out6 <prefix-list-name_str>
    set remove-private-as6 {enable | disable}
    set route-map-in6 <routemap-name_str>
    set route-map-out6 <routemap-name_str>
    set route-reflector-client6 {enable | disable}
    set route-server-client6 {enable | disable}
    set send-community6 {both | disable | extended | standard}
    set soft-reconfiguration6 {enable | disable}
    set unsuppress-map6 <route-map-name_str>
config network6

```

```
config redistribute6
end
```

RIPng — RIP and IPv6

RIP next generation, or RIPng, is the version of RIP that supports IPv6.

This is an example of a typical small network configuration using RIPng routing.

Your internal R&D network is working on a project for a large international telecom company that uses IPv6. For this reason, you have to run IPv6 on your internal network and you have decided to use only IPv6 addresses.

Your network has two FortiGate units running the RIPng dynamic routing protocol. Both FortiGate units are connected to the ISP router and the internal network. This configuration provides some redundancy for the R&D internal network enabling it to reach the internet at all times.

This section includes the following topics:

- [Network layout and assumptions](#)
- [Configuring the FortiGate units system information](#)
- [Configuring RIPng on FortiGate units](#)
- [Configuring other network devices](#)
- [Testing the configuration](#)
- [Debugging IPv6 on RIPng](#)

Network layout and assumptions

Basic network layout

Your internal R&D network is working on a project for a large international telecom company that uses IPv6. For this reason, you have to run IPv6 on your internal network and you have decided to use only IPv6 addresses.

Your network has two FortiGate units running the RIPng dynamic routing protocol. Both FortiGate units are connected to the ISP router and the internal network. This configuration provides some redundancy for the R&D internal network enabling it to reach the internet at all times.

All internal computers use RIP routing, so no static routing is required. And all internal computers use IPv6 addresses.


Where possible in this example, the default values will be used or the most general settings. This is intended to provide an easier configuration that will require less troubleshooting.

In this example the routers, networks, interfaces used, and IP addresses are as follows:

Rip example network topology

Network	Router	Interface & Alias	IPv6 address
R&D	Router1	port1 (internal)	2002:A0B:6565:0:0:0:0:0
		port2 (ISP)	2002:AC14:7865:0:0:0:0:0
	Router2	port1 (internal)	2002:A0B:6566:0:0:0:0:0
		port2 (ISP)	2002:AC14:7866:0:0:0:0:0

Network topology for the IPv6 RIPng example



Assumptions

The following assumptions have been made concerning this example:

- All FortiGate units have 5.0+ firmware, and are running factory default settings.
- All CLI and web-based manager navigation assumes the unit is running in NAT/Route operating mode, with VDOMs disabled.
- All FortiGate units have interfaces labeled port1 and port2 as required.
- All firewalls have been configured for each FortiGate unit to allow the required traffic to flow across interfaces.
- All network devices are support IPv6 and are running RIPng.

Configuring the FortiGate units system information

Each FortiGate unit needs IPv6 enabled, a new hostname, and interfaces configured.

To configure system information on Router1 - web-based manager:

1. Go to **Dashboard**.
2. For **Host name**, select **Change**.
3. Enter "Router1".
4. Go to **System > Feature Visibility**.
5. Enable **IPv6** and click **Apply**.
6. Go to **Network > Interfaces**.
7. Edit port1 (internal) interface.
8. Set the following information, and select **OK**.

Alias	internal
IP/Netmask	2002:A0B:6565::/0
Administrative Access	HTTPS SSH PING
Description	Internal RnD network
Administrative Status	Up

9. Edit port2 (ISP) interface.
10. Set the following information, and select **OK**.

Alias	ISP
IP/Netmask	2002:AC14:7865::/0
Administrative Access	HTTPS SSH PING
Description	ISP and internet
Administrative Status	Up

To configure system information on Router1 - CLI:

```

config system global
    set hostname Router1
    set gui-ipv6 enable
end
config system interface
    edit port1
        set alias internal
        set allowaccess https ping ssh
        set description "Internal RnD network"
        config ipv6
            set ip6-address 2002:a0b:6565::/0
        end
    end
    next
    edit port2
        set alias ISP
        set allowaccess https ping ssh
        set description "ISP and internet"
        config ipv6
            set ip6-address 2002:AC14:7865::
        end
    end
end

```

To configure system information on Router2 - web-based manager:

1. Go to **Dashboard**.
2. For **Host name**, select **Change**.
3. Enter "Router2".
4. Go to **System > Feature Visibility**.
5. Enable **IPv6** and click **Apply**.
6. Go to **Network > Interfaces**.
7. Edit port1 (internal) interface.
8. Set the following information, and select **OK**.

Alias	internal
IP/Netmask	2002:A0B:6566::/0

Administrative Access	HTTPS SSH PING
Description	Internal RnD network
Administrative Status	Up

9. Edit port2 (ISP) interface.
10. Set the following information, and select **OK**.

Alias	ISP
IP/Netmask	2002:AC14:7866::/0
Administrative Access	HTTPS SSH PING
Description	ISP and internet
Administrative Status	Up

To configure system information on Router2 - CLI:

```

config system global
    set hostname Router2
    set gui-ipv6 enable
end
config system interface
    edit port1
        set alias internal
        set allowaccess https ping ssh
        set description "Internal RnD network"
        config ipv6
            set ip6-address 2002:a0b:6566::/0
        end
    end
    next
    edit port2
        set alias ISP
        set allowaccess https ping ssh
        set description "ISP and internet"
        config ipv6
            set ip6-address 2002:AC14:7866::
        end
    end
end

```

Configuring RIPng on FortiGate units

Now that the interfaces are configured, you can configure RIPng on the FortiGate units.

There are only two networks and two interfaces to include — the internal network, and the ISP network. There is no redistribution, and no authentication. In RIPng there is no specific command to include a subnet in the RIP broadcasts. There is also no information required for the interfaces beyond including their name.

As this is a CLI only configuration, configure the ISP router and the other FortiGate unit as neighbors. This was not part of the previous example as this feature is not offered in the web-based manager. Declaring neighbors in the configuration like this will reduce the discovery traffic when the routers start up.

Since RIPng is not supported in the web-based manager, this section will only be entered in the CLI.

To configure RIPng on Router1 - CLI:

```
config router ripng
config interface
edit port1
next
edit port2
end
config neighbor
edit 1
set interface port1
set ipv6 2002:a0b:6566::/0
next
edit 2
set interface port2
set ipv6 2002:AC14:7805::/0
end
```

To configure RIPng on Router2 - CLI:

```
config router ripng
config interface
edit port1
next
edit port2
end
config neighbor
edit 1
set interface port1
set ipv6 2002:a0b:6565::/0
next
edit 2
set interface port2
set ipv6 2002:AC14:7805::/0
end
```

Configuring other network devices

The other devices on the internal network all support IPv6, and are running RIPng where applicable. They only need to know the internal interface network addresses of the FortiGate units.

The ISP routers need to know the FortiGate unit information such as IPv6 addresses.

Testing the configuration

In addition to normal testing of your network configuration, you must also test the IPv6 part of this example.

For troubleshooting problems with your network, see the [FortiOS Handbook Troubleshooting](#) chapter.

Testing the IPv6 RIPng information

There are some commands to use when checking that your RIPng information is correct on your network. These are useful to check on your RIPng FortiGate units on your network. Comparing the output between devices will help you understand your network better, and also track down any problems.

```
diagnose ipv6 address list
```

View the local scope IPv6 addresses used as next-hops by RIPng on the FortiGate unit.

```
diagnose ipv6 route list
```

View ipv6 addresses that are installed in the routing table.

```
get router info6 routing-table
```

View the routing table. This information is almost the same as the previous command (`diagnose ipv6 route list`) however it is presented in an easier to read format.

```
get router info6 rip interface external
```

View brief output on the RIP information for the interface listed. The information includes if the interface is up or down, what routing protocol is being used, and whether passive interface or split horizon are enabled.

```
get router info6 neighbor-cache list
```

View the IPv6/MAC address mapping. This also displays the interface index and name associated with the address.

Debugging IPv6 on RIPng

The debug commands are very useful to see what is happening on the network at the packet level. There are a few changes to debugging the packet flow when debugging IPv6.

The following CLI commands specify both IPv6 and RIP, so only RIPng packets will be reported. The output from these commands will show you the RIPng traffic on your FortiGate unit including RECV, SEND, and UPDATE actions.

The addresses are in IPv6 format.

```
diagnose debug enable
diagnose ipv6 router rip level info
diagnose ipv6 router rip all enable
```

These three commands will:

Turn on debugging in general

Set the debug level to information, a verbose reporting level

Turn on all RIP router settings

Part of the information displayed from the debugging is the metric (hop count). If the metric is 16, then that destination is unreachable since the maximum hop count is 15.

In general, you should see an update announcement, followed by the routing table being sent out, and a received reply in response.

IPv6 RSSO support

RADIUS Single Sign-On (RSSO) is supported in IPv6, but can only be configured in the CLI:

```
config firewall policy6
  edit <id>
    set rso enable
    set fall-through-unauthenticated enable
  end
```

IPv6 IPS

IPv6 IPS signature scan can be enabled by interface policy. The user can create an normal IPS sensor and assign it to the IPv6 interface policy.

```
config firewall interface-policy6
  edit 1
    set interface "port1"
    set srcaddr6 "all"
    set dstaddr6 "all"
    set service6 "ANY"
    set ips-sensor-status enable
    set ips-sensor "all_default"
  next
end
```

Blocking IPv6 packets by extension headers

FortiOS can now block IPv6 packets based on the extension headers, using the CLI syntax:

```
config firewall ipv6-eh-filter.
```

The following commands are now available:

- `set hop-opt {disable | enable}`: Block packets with Hop-by-Hop Options header.
- `set dest-opt {disable | enable}`: Block packets with Destination Options header.
- `set hdopt-type <integer>`: Block specific Hop-by-Hop and/or Destination Option types (maximum 7 types, each between 0 and 255).
- `set routing {disable | enable}`: Block packets with Routing header.
- `set routing-type <integer>`: Block specific Routing header types (maximum 7 types, each between 0 and 255).
- `set fragment {disable | enable}`: Block packets with Fragment header.
- `set auth {disable | enable}`: Block packets with Authentication header.
- `set no-next {disable | enable}`: Block packets with No Next header.

IPv6 denial of service policies

Denial of Service (DoS) policies can now be configured by going to **Policy & Objects > IPv6 DoS Policy**. For more information, refer to the “Interface Policies” section of the [FortiOS Handbook Firewall](#) chapter.

Configure hosts in an SNMP v1/2c community to send queries or receive traps

When you add a host to an SNMP v1/2c community you can now decide whether the FortiGate unit will accept queries from the host or whether the FortiGate unit will send traps to the host. You can also configure the host for both traps and queries. You can add up to 16 IPv4 hosts and up to 16 IPv6 hosts.

Use the following command to add two hosts to an SNMP community:

```
config system snmp community
  config hosts
    edit 1
      set interface port1
      set ip 172.20.120.1
      set host-type query
    end
  config hosts6
```

```

edit 1
    set interface port6
    set ip 2001:db8:0:2::30
    set host-type trap
end

```

IPv6 PIM sparse mode multicast routing

FortiOS supports PIM sparse mode multicast routing for IPv6 multicast (multicast6) traffic and is compliant with [RFC 4601](#). You can use the following command to configure IPv6 PIM sparse multicast routing.

```

config router multicast6
    set multicast-routing {enable | disable}
    config interface
        edit <interface-name>
            set hello-interval <1-65535 seconds>
            set hello-holdtime <1-65535 seconds>
        end
    config pim-sm-global
    config rp-address
        edit <index>
            set ipv6-address <ipv6-address>
        end
    end
end

```

The following diagnose commands for IPv6 PIM sparse mode are also available:

```

diagnose ipv6 multicast status
diagnose ipv6 multicast vif
diagnose ipv6 multicast mroute

```

IPv6 Neighbor Discovery Proxy

The following is an example configuration of a FortiGate using ND Proxy. Some of these configuration steps have been covered elsewhere, but are shown here to demonstrate how they all work together to achieve the desired effect.

Steps:

- Create zone for ND proxy use that includes the upstream and downstream interfaces.
- Create policies to allow ICMPv6 and DHCPv6 traffic.
- Enable ND Proxy on the interfaces.
- Enable "autoconf" on the upstream interface.

1. Add a zone including wan and lan.

It is possible to use firewall and multicast policies that don't use a zone, but using a zone simplifies the configuration, especially if you have more than two interfaces.

```

config system zone
    edit ndproxy_zone
        set interface wan lan
    end

```



On some models the "lan" interface is named "internal".

2. Add forward firewall policy and multicast policy to allow at least ICMPv6 and DHCPv6 traffic.

```
config firewall multicast-policy6
edit 0
    set srcintf ndproxy_zone
    set dstintf ndproxy_zone
    set srcaddr all
    set dstaddr all
end
```

and

```
config firewall policy6
edit 0
    set srcintf ndproxy_zone
    set dstintf ndproxy_zone
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ALL
end
```

3. Enable ND proxy on WAN and LAN.

```
config system nd-proxy
    set status enable
    set member wan lan
end
```

4. Enable `autoconf` on the upstream interface.

RA received on the other interface(s) will be dropped.

```
config system interface
edit wan
...
config ipv6
    set autoconf enable
end
end
```

Network defense

This section describes in general terms the means by which attackers can attempt to compromise your network using attacks at the network level rather than through application vulnerabilities, and steps you can take to protect it. The goal of an attack can be as complex as gaining access to your network and the privileged information it contains, or as simple as preventing customers from accessing your web server.

Because of popular media, many people are aware of viruses and other malware as a threat against their computers and data, but some of the most costly malicious attack in history have been against networks. A 2016 study found that a single DDoS attack could cost a company over \$1.6 million. Depending on the size and type of company the areas of expense can be:

- Changes in credit and insurance ratings
- Overtime payment to employees
- Hiring new employees to increase IT staff
- PR expenses to restore a company's reputation
- Upgrading infrastructure and software
- Customer compensation

The following topics are included in this section:

- [Monitoring](#)
- [Blocking external probes](#)
- [Defending against DoS attacks](#)

Monitoring

Monitoring, in the form of logging, alert email, and SNMP, does not directly protect your network. But monitoring allows you to review the progress of an attack, whether afterwards or while in progress. How the attack unfolds may reveal weaknesses in your preparations. The packet archive and sniffer policy logs can reveal more details about the attack. Depending on the detail in your logs, you may be able to determine the attacker's location and identity.

While log information is valuable, you must balance the log information with the resources required to collect and store it.

Blocking external probes

Protection against attacks is important, but attackers often use vulnerabilities and network tools to gather information about your network to plan an attack. It is often easier to prevent an attacker from learning important details about your network than to defend against an attack designed to exploit your particular network.

Attacks are often tailored to the hardware or operating system of the target, so reconnaissance is often the first step. The IP addresses of the hosts, the open ports, and the operating systems the hosts are running is invaluable information to an attacker. Probing your network can be as simple as an attacker performing an address sweep or port scan to a more involved operation like sending TCP packets with invalid combinations of flags to see how your firewall reacts.

Address sweeps

An address sweep is a basic network scanning technique to determine which addresses in an address range have active hosts. A typical address sweep involves sending an ICMP ECHO request (a ping) to each address in an address range to attempt to get a response. A response signifies that there is a host at this address that responded to the ping. It then becomes a target for more detailed and potentially invasive attacks.

Address sweeps do not always reveal all the hosts in an address range because some systems may be configured to ignore ECHO requests and not respond, and some firewalls and gateways may be configured to prevent ECHO requests from being transmitted to the destination network. Despite this shortcoming, Address sweeps are still used because they are simple to perform with software tools that automate the process.

Use the `icmp_sweep` anomaly in a DoS policy to protect against address sweeps.

There are a number of IPS signatures to detect the use of ICMP probes that can gather information about your network. These signatures include `AddressMask`, `Traceroute`, `ICMP.Invalid.Packet.Size`, and `ICMP.Oversized.Packet`. Include ICMP protocol signatures in your IPS sensors to protect against these probes/attacks.

Port scans

Potential attackers may run a port scan on one or more of your hosts. This involves trying to establish a communication session to each port on a host. If the connection is successful, a service may be available that the attacker can exploit.

Use the DoS anomaly check for `tcp_port_scan` to limit the number of sessions (complete and incomplete) from a single source IP address to the configured threshold. If the number of sessions exceed the threshold, the configured action is taken.

Use the DoS anomaly check for `udp_scan` to limit UDP sessions in the same way.

Probes using IP traffic options

Every TCP packet has space reserved for eight flags or control bits. They are used for communicating various control messages. Although space in the packet is reserved for all eight, there are various combinations of flags that should never happen in normal network operation. For example, the SYN flag, used to initiate a session, and the FIN flag, used to end a session, should never be set in the same packet.

Attackers may create packets with these invalid combinations to test how a host will react. Various operating systems and hardware react in different ways, giving a potential attackers clues about the components of your network.

The IPS signature `TCP.Bad.Flags` detects these invalid combinations. The default action is pass though you can override the default and set it to **Block** in your IPS sensor.

Configure packet replay and TCP sequence checking

The `anti-replay` command in the CLI allows you to set the level of checking for packet replay and TCP sequence checking (or TCP Sequence (SEQ) number checking). All TCP packets contain a Sequence Number (SEQ) and an Acknowledgment Number (ACK). The TCP protocol uses these numbers for error free end-to-end communications. TCP sequence checking can also be used to validate individual packets.

FortiGate units use TCP sequence checking to make sure that a packet is part of a TCP session. By default, if a packet is received with sequence numbers that fall out of the expected range, the FortiGate unit drops the packet. This is normally a desired behavior, since it means that the packet is invalid. But in some cases you may

want to configure different levels of `anti-replay` checking if some of your network equipment uses non-RFC methods when sending packets.

Configure the anti-replay CLI command:

```
config system global
    set anti-replay {disable | loose | strict}
end
```

You can set `anti-replay` protection to the following settings:

- `disable` — No anti-replay protection.
- `loose` — Perform packet sequence checking and ICMP anti-replay checking with the following criteria:
 - The SYN, FIN, and RST bit can not appear in the same packet.
 - The FortiGate unit does not allow more than one ICMP error packet through before it receives a normal TCP or UDP packet.
 - If the FortiGate unit receives an RST packet, and `check-reset-range` is set to `strict`, the FortiGate unit checks to determine if its sequence number in the RST is within the un-ACKed data and drops the packet if the sequence number is incorrect.
- `strict` — Performs all of the loose checking but for each new session also checks to determine if the TCP sequence number in a SYN packet has been calculated correctly and started from the correct value for each new session. Strict anti-replay checking can also help prevent SYN flooding.

If any packet fails a check it is dropped.



The anti-replay setting only affects non-accelerated traffic.

Configure ICMP error message verification

Enable ICMP error message verification to ensure an attacker can not send an invalid ICMP error message.

```
config system global
    check-reset-range {disable | strict}
end
```

- `disable` — the FortiGate unit does not validate ICMP error messages.
- `strict` — enable ICMP error message checking.

If the FortiGate unit receives an ICMP error packet that contains an embedded IP(A,B) | TCP(C,D) header, then if FortiOS can locate the A:C->B:D session it checks to make sure that the sequence number in the TCP header is within the range recorded in the session. If the sequence number is not in range then the ICMP packet is dropped. Strict checking also affects how the anti-replay option checks packets.

Protocol header checking

Select the level of checking performed on protocol headers.

```
config system global
    check-protocol-header {loose | strict}
end
```

- `loose` — the FortiGate unit performs basic header checking to verify that a packet is part of a session and should be processed. Basic header checking includes verifying that the layer-4 protocol header length, the IP header length, the IP version, the IP checksum, IP options are correct, etc.

- `strict` — the FortiGate unit does the same checking as above plus it verifies that ESP packets have the correct sequence number, SPI, and data length.

If the packet fails header checking it is dropped by the FortiGate unit.

Evasion techniques

Attackers employ a wide range of tactics to try to disguise their techniques. If an attacker disguises a known attack in such a way that it is not recognized, the attack will evade your security and possibly succeed. FortiGate security recognizes a wide variety of evasion techniques and normalizes data traffic before inspecting it.

Packet fragmentation

Information sent across local networks and the Internet is encapsulated in packets. There is a maximum allowable size for packets and this maximum size varies depending on network configuration and equipment limitations. If a packet arrives at a switch or gateway and it is too large, the data it carries is divided among two or more smaller packets before being forwarded. This is called fragmentation.

When fragmented packets arrive at their destination, they are reassembled and read. If the fragments do not arrive together, they must be held until all of the fragments arrive. Reassembly of a packet requires all of the fragments.

The FortiGate unit automatically reassembles fragmented packets before processing them because fragmented packets can evade security measures. This reassembly of packets affects TCP, UDP and IP packets. There can be some variation though in what process does the reassembling. The IPS engine, nTurbo and the kernel all can do defragmentation.

For example, you have configured the FortiGate unit to block access to the example.org web site. Any checks for example.com will fail if a fragmented packet arrives and one fragment contains `http://www.exa` while the other contains `mple.com/`. Viruses and malware can be fragmented and avoid detection in the same way. The FortiGate unit will reassemble fragmented packets before examining network data to ensure that inadvertent or deliberate packet fragmentation does not hide threats in network traffic.

Non-standard ports

Most traffic is sent on a standard port based on the traffic type. The FortiGate unit recognizes most traffic by packet content rather than the TCP/UDP port and uses the proper IPS signatures to examine it. Protocols recognized regardless of port include DHCP, DNP3, FTP, HTTP, IMAP, MS RPC, NNTP, POP3, RSTP, SIP, SMTP, and SSL, as well as the supported IM/P2P application protocols.

In this way, the FortiGate unit will recognize HTTP traffic being sent on port 25 as HTTP rather than SMTP, for example. Because the protocol is correctly identified, the FortiGate unit will examine the traffic for any enabled HTTP signatures.

Negotiation codes

Telnet and FTP servers and clients support the use of negotiation information to allow the server to report what features it supports. This information has been used to exploit vulnerable servers. To avoid this problem, the FortiGate unit removes negotiation codes before IPS inspection.

HTTP URL obfuscation

Attackers encode HTML links using various formats to evade detection and bypass security measures. For example, the URL `www.example.com/cgi.bin` could be encoded in a number of ways to avoid detection but still

work properly, and be interpreted the same, in a web browser.

The FortiGate prevents the obfuscation by converting the URL to ASCII before inspection.

HTTP URL obfuscation types

Encoding type	Example
No encoding	http://www.example.com/cgi.bin/
Decimal encoding	http://www.example.com/cgi.bin/
URL encoding	http://www.example.com/%43%47%49%2E%42%49%4E%2F
ANSI encoding	http://www.example.com/%u0063%u0067%u0069%u002E%u0062%u0069%u006E/
Directory traversal	http://www.example.com/cgi.bin/test/..

HTTP header obfuscation

The headers of HTTP requests or responses can be modified to make the discovery of patterns and attacks more difficult. To prevent this, the FortiGate unit will:

- remove junk header lines
- reassemble an HTTP header that's been folded onto multiple lines
- move request parameters to HTTP POST body from the URL

The message is scanned for any enabled HTTP IPS signatures once these problems are corrected.

HTTP body obfuscation

The body content of HTTP traffic can be hidden in an attempt to circumvent security scanning. HTTP content can be GZipped or deflated to prevent security inspection. The FortiGate unit will uncompress the traffic before inspecting it.

Another way to hide the contents of HTTP traffic is to send the HTTP body in small pieces, splitting signature matches across two separate pieces of the HTTP body. The FortiGate unit reassembles these 'chunked bodies' before inspection.

Microsoft RPC evasion

Because of its complexity, the Microsoft Remote Procedure Call protocol suite is subject to a number of known evasion techniques, including:

- SMB-level fragmentation
- DCERPC-level fragmentation
- DCERPC multi-part fragmentation

- DCERPC UDP fragmentation
- Multiple DCERPC fragments in one packet

The FortiGate unit reassembles the fragments into their original form before inspection.

Defending against DoS attacks

A denial of service is the result of an attacker sending an abnormally large amount of network traffic to a target system. Having to deal with the traffic flood slows down or disables the target system so that legitimate users can not use it for the duration of the attack.

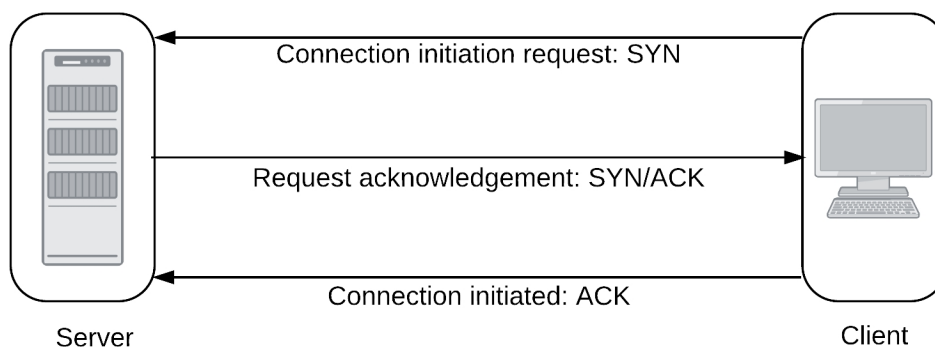
Any network traffic the target system receives has to be examined, and then accepted or rejected. TCP, UDP, and ICMP traffic is most commonly used, but a particular type of TCP traffic is the most effective. TCP packets with the SYN flag are the most efficient DoS attack tool because of how communication sessions are started between systems.

The “three-way handshake”

Communication sessions between systems start with establishing a TCP/IP connection. This is a simple three step process, sometimes called a “three-way handshake,” initiated by the client attempting to open the connection.

1. The client sends a TCP packet with the SYN flag set. With the SYN packet, the client informs the server of its intention to establish a connection.
2. If the server is able to accept the connection to the client, it sends a packet with the SYN and the ACK flags set. This simultaneously acknowledges the SYN packet the server has received, and informs the client that the server intends to establish a connection.
3. To acknowledge receipt of the packet and establish the connection, the client sends an ACK packet.

Establishing a TCP/IP connection



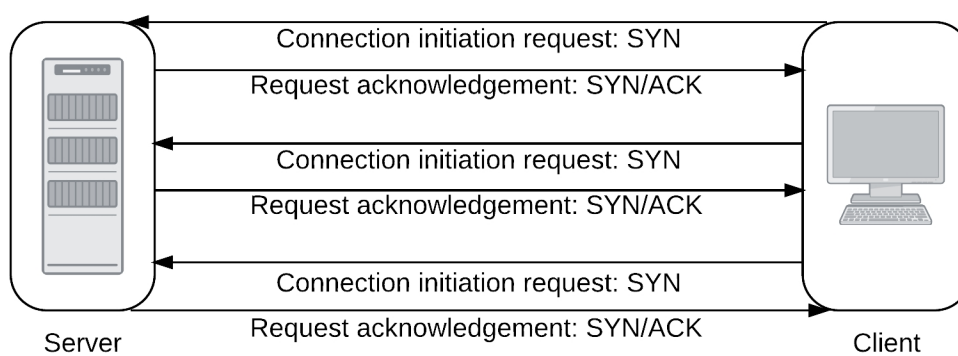
The three-way handshake is a simple way for the server and client to each agree to establish a connection and acknowledge the other party expressing its intent. Unfortunately, the three-way handshake can be used to interfere with communication rather than facilitate it.

SYN flood

When a client sends a SYN packet to a server, the server creates an entry in its session table to keep track of the connection. The server then sends a SYN+ACK packet expecting an ACK reply and the establishment of a connection.

An attacker intending to disrupt a server with a denial of service (DoS) attack can send a flood of SYN packets and not respond to the SYN+ACK packets the server sends in response. Networks can be slow and packets can get lost so the server will continue to send SYN+ACK packets until it gives up, and removes the failed session from the session table. If an attacker sends enough SYN packets to the server, the session table will fill completely, and further connection attempts will be denied until the incomplete sessions time out. Until this happens, the server is unavailable to service legitimate connection requests.

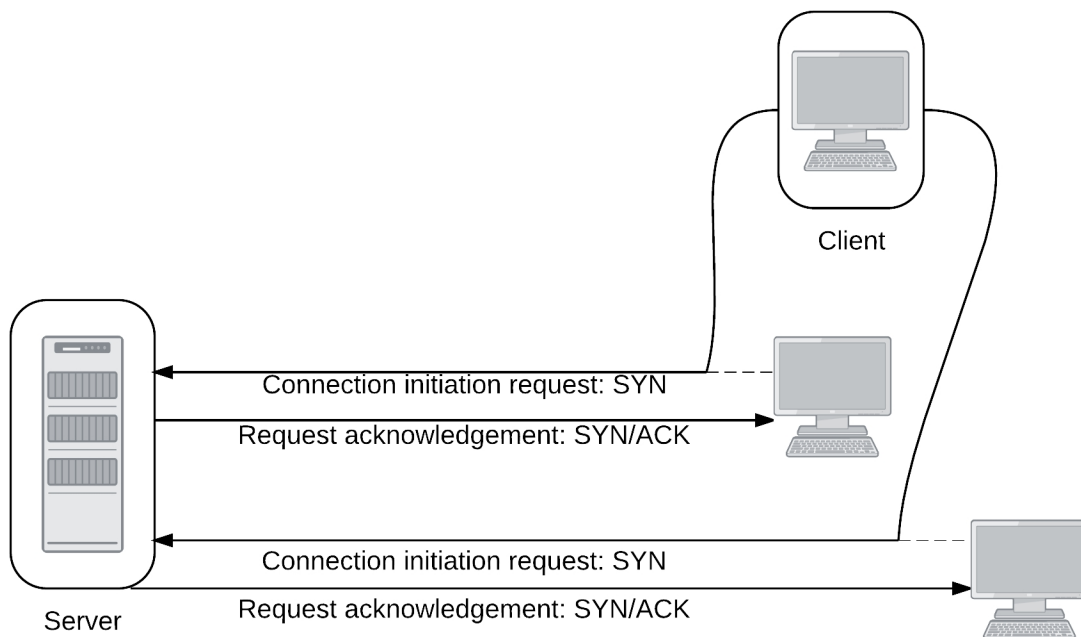
A single client launches a SYN flood attack



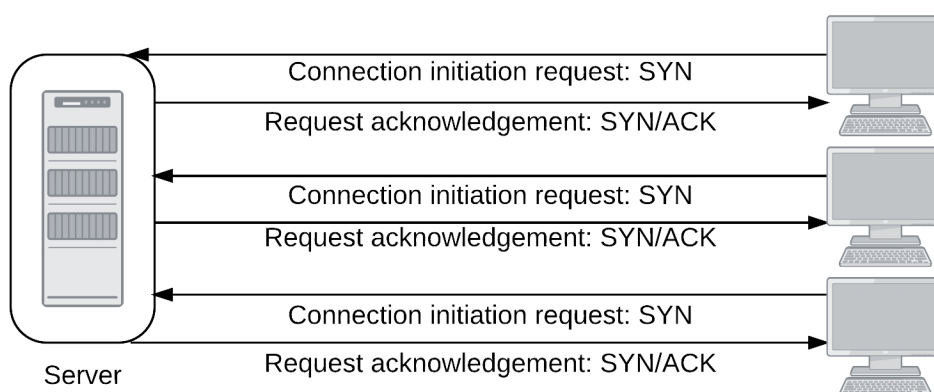
SYN floods are seldom launched from a single address so limiting the number of connection attempts from a single IP address is not usually effective.

SYN spoofing

With a flood of SYN packets coming from a single attacker, you can limit the number of connection attempts from the source IP address or block the attacker entirely. To prevent this simple defense from working, or to disguise the source of the attack, the attacker may spoof the source address and use a number of IP addresses to give the appearance of a distributed denial of service (DDoS) attack. When the server receives the spoofed SYN packets, the SYN+ACK replies will go to the spoofed source IP addresses which will either be invalid, or the system receiving the reply will not know what to do with it.

A client launches a SYN spoof attack**DDoS SYN flood**

The most severe form of SYN attack is the distributed SYN flood, one variety of distributed denial of service attack (DDoS). Like the SYN flood, the target receives a flood of SYN packets and the ACK+SYN replies are never answered. The attack is distributed across multiple sources sending SYN packets in a coordinated attack.

Multiple attackers launch a distributed SYN flood

The distributed SYN flood is more difficult to defend against because multiple clients are capable of creating a larger volume of SYN packets than a single client. Even if the server can cope, the volume of traffic may

overwhelm a point in the network upstream of the targeted server. The only defense against this is more bandwidth to prevent any choke points.

Configuring the SYN threshold to prevent SYN floods

The preferred primary defense against any type of SYN flood is the DoS anomaly check for `tcp_syn_flood` threshold. The threshold value sets an upper limit on the number of new incomplete TCP connections allowed per second. If the number of incomplete connections exceeds the threshold value, and the action is set to **Pass**, the FortiGate unit will allow the SYN packets that exceed the threshold. If the action is set to **Block**, the FortiGate unit will block the SYN packets that exceed the threshold, but it will allow SYN packets from clients that send another SYN packet.

The tools attackers use to generate network traffic will not send a second SYN packet when a SYN+ACK response is not received from the server. These tools will not “retry.” Legitimate clients will retry when no response is received, and these retries are allowed even if they exceed the threshold with the action set to **Block**.

SYN proxy

FortiGate units with network acceleration hardware, whether built-in or installed in the form of an add-on module, offer a third action for the `tcp_syn_flood` threshold. Instead of **Block** and **Pass**, you can choose to **Proxy** the incomplete connections that exceed the threshold value.

When the `tcp_syn_flood` threshold action is set to **f**, incomplete TCP connections are allowed as normal as long as the configured threshold is not exceeded. If the threshold is exceeded, the FortiGate unit will intercept incoming SYN packets from clients and respond with a SYN+ACK packet. If the FortiGate unit receives an ACK response as expected, it will “replay” this exchange to the server to establish a communication session between the client and the server, and allow the communication to proceed.

Other flood types

UDP and ICMP packets can also be used for DoS attacks, though they are less common. TCP SYN packets are so effective because the target receives them and maintains a session table entry for each until they time out. Attacks using UDP or ICMP packets do not require the same level of attention from a target, rendering them less effective. The target will usually drop the offending packets immediately, closing the session.

Use the `udp_flood` and `icmp_flood` thresholds to defend against these DoS attacks.

DoS policies

DDoS attacks vary in nature and intensity. Attacks aimed at saturating the available bandwidth upstream of your service can only be countered by adding more bandwidth. DoS policies can help protect against DDoS attacks that aim to overwhelm your server resources.

DoS policy recommendations

- Use and configure DoS policies to appropriate levels based on your network traffic and topology. This will help drop traffic if an abnormal amount is received.
- It is important to set a good threshold. The threshold defines the maximum number of sessions/packets per second of normal traffic. If the threshold is exceeded, the action is triggered. Threshold defaults are general recommendations, although your network may require very different values.
- One way to find the correct values for your environment is to set the action to **Pass** and enable logging. Observe the logs and adjust the threshold values until you can determine the value at which normal traffic begins to generate attack reports. Set the threshold above this value with the margin you want. Note that the smaller the margin, the


more protected your system will be from DoS attacks, but your system will also be more likely to generate false alarms.

Inside FortiOS: Denial of Service (DoS) protection

FortiOS DoS protection maintains network integrity and performance by identifying and blocking harmful IPv4 and IPv6-based denial of service (DoS) attacks.

About DoS and DDoS attacks

A denial of service (DoS) occurs when an attacker overwhelms server resources by flooding a target system with anomalous data packets, rendering it unable to service genuine users. A distributed denial of service (DDoS) occurs when an attacker uses a master computer to control a network of compromised systems, otherwise known as a 'botnet', which collectively inundates the target system with excessive anomalous data packets.

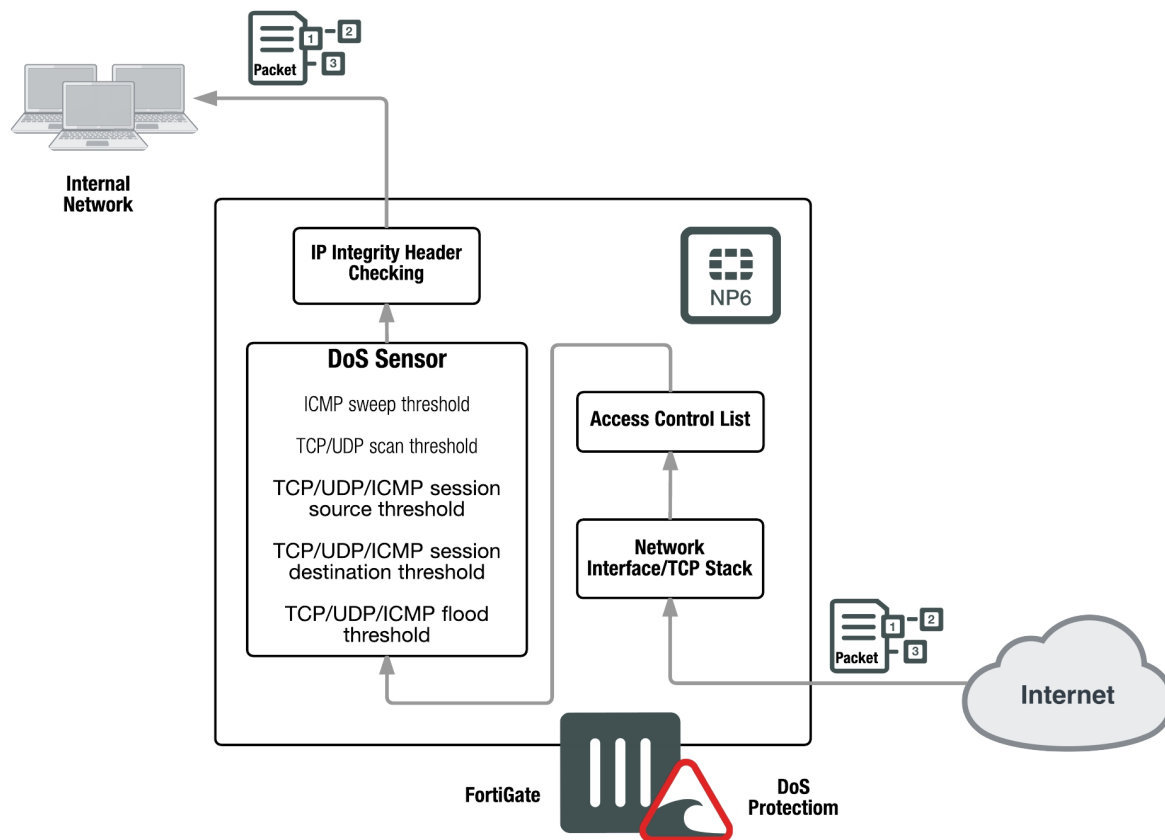


FortiOS DoS and DDoS protection

FortiOS DoS protection identifies potentially harmful traffic that could be part of a DoS or a DDoS attack by looking for specific traffic anomalies. Traffic anomalies that become DoS attacks include: TCP SYN floods, UDP floods, ICMP floods, TCP port scans, TCP session attacks, UDP session attacks, ICMP session attacks, and ICMP sweep attacks. Only traffic identified as part of a DoS attack is blocked; connections from legitimate users are processed normally.

FortiOS applies DoS protection very early in its traffic processing sequence to minimize the effect of a DoS attack on FortiOS system performance. DoS protection is the first step for packets after they are received by a FortiGate interface. Potential DoS attacks are detected and blocked before the packets are sent to other FortiOS systems.

FortiOS also includes an access control list feature that is implemented next. This accelerated ACL technology uses NP6 processors to block traffic (including DoS attacks) by source and destination address and service again before the packets are sent to the FortiGate CPU.



FortiOS DoS protection can operate in a standard configuration or operate out of band in sniffer mode, also known as one-arm mode, similar to intrusion detection systems. When operating in sniffer mode the FortiGate unit detects attacks and logs them without blocking them.

FortiOS DoS policies determine the course of action to take when anomalous traffic reaches a configured packet rate threshold. You can block an attacker, block an interface, block an attacker and interface, or allow traffic to pass through for monitoring purposes. This allows you to maintain network security by gathering information about attacks, monitor potentially offending traffic, or block offenders for the most protection.

FortiGates with NP6 processors also support synproxy DoS protection. An NP6-accelerated TCP SYN proxy offloads the three-way TCP handshake TCP SYN anomaly checking DoS protection to NP6 processors.

FortiOS DDoS prevention

In addition to using DoS protection for protection against DoS attacks, FortiOS includes a number of features that prevent the spread of Botnet and C&C activity. Mobile Malware or Botnet and C&C protection keeps Botnet and C&C code from entering a protected network and compromising protected systems. As a result, systems on the protected network cannot become Botnet clients.


In addition, FortiOS can monitor and block outgoing Botnet connection attempts. Monitoring allows you to find and remove Botnet clients from your network and blocking prevents infected systems from communicating with Botnet sites.

Configuration options

Choose the standard configuration for maximum protection or configure sniffer mode to gather information.

Standard configuration

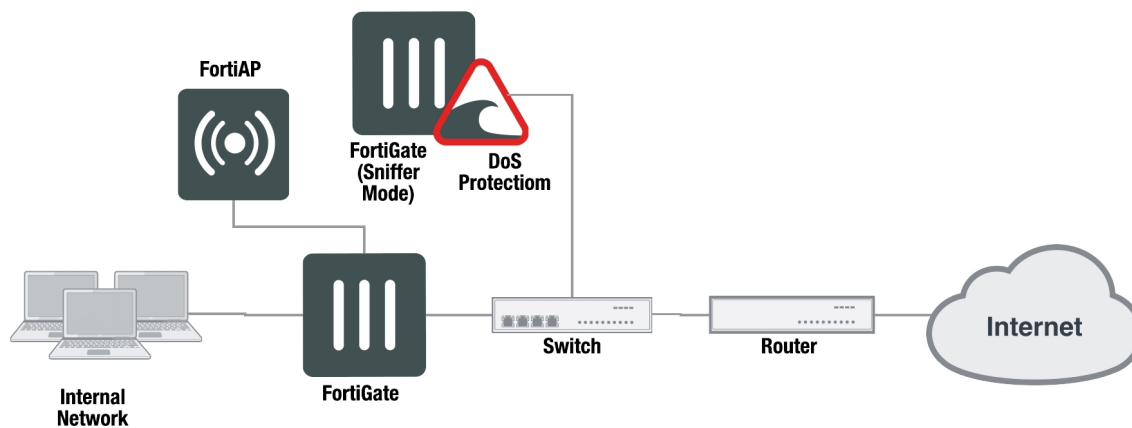
DoS protection is commonly configured on a FortiGate unit that connects a private or DMZ network to the Internet or on a FortiWiFi unit that connects a wireless LAN to an internal network and to the Internet. All Internet traffic or wireless LAN traffic passes through DoS protection in the FortiGate unit or the FortiWiFi unit.



Out of band configuration (sniffer mode)

A FortiGate unit in sniffer mode operates out of band as a one-armed Intrusion Detection System by detecting and reporting attacks. It does not process network traffic nor does it take action against threats. The FortiGate interface operating in sniffer mode is connected to a Test Access Point (TAP) or a Switch Port Analyzer (SPAN) port that processes all of the traffic to be analyzed. The TAP or SPAN sends a copy of the switch traffic to the out of band FortiGate for analysis.

FortiOS records log messages and sends alerts to system administrators when a DoS attack is detected. IDS scanning does not affect network performance or network traffic if the IDS fails or goes offline.



DoS policies

DoS policies provide effective and early DoS detection while remaining light on system resources. They are configured to monitor and to stop traffic with abnormal patterns or attributes. The DoS policy recognizes traffic as a threat when the traffic reaches a user-configured packet rate threshold. The policy then determines the appropriate action. In addition to choosing whether or not to log each type of anomaly, you can choose to pass or block threats.

DoS policy anomaly protection is applied to all incoming traffic to a single FortiGate interface, but you can narrow policies by specifying service, source address, and destination address. The FortiGate unit processes DoS policies in their own respective order first, followed by all other firewall policies.

Hardware acceleration

Hardware acceleration enhances protection and increases the efficiency of your network. FortiOS integrated Content Processors (CPs), Network Processors (NPs), and Security Processors (SPs) accelerate specialized security processing. DoS SYN proxy protection is built in to NP6 processors and many Fortinet Security Processors, like the CE4, XE2, and FE8, to guard against TCP SYN floods. TCP packets with the SYN flag are the most efficient DoS attack tool because of how communication sessions are initiated between systems. NP6 and SP processors can offload TCP SYN flood attack detection and blocking. The SP module increases a FortiGate unit's capacity to protect against TCP SYN flood attacks while minimizing the effect of attacks on the FortiGate unit's overall performance and the network performance. The result is improved capacity and overall system performance.

The FortiGuard Center

The FortiGuard Center shows information on all the most recent FortiGuard news, including information concerning zero-day research and hot intrusion detections. Research papers are also available that concern a variety of current security issues.

To view recent developments, go to <http://www.fortiguard.com/static/intrusionprevention.html>.

Policy configuration

The firewall policies of the FortiGate are one of the most important aspects of the appliance. There are a lot of building blocks and configurations involved in setting up a firewall and it within the policies that a lot of these components come together to form a cohesive unit to perform the firewall's main function, analyzing network traffic and responding appropriately to the results of that analysis.

There are a few different kinds of policies and in most cases these are further divided into IPv4 and IPv6 versions:

- [IPv4 policy](#) - used for managing traffic going through the appliance using IPv4 protocols
- [IPv6 policy](#) - used for managing traffic going through the appliance using IPv6 protocols
- [NAT64 policy](#) - used for managing traffic going through the appliance that converts from IPv6 on the incoming interface to IPv4 on the outgoing interface
- [NAT46 policy](#) - used for managing traffic going through the appliance that converts from IPv4 on the incoming interface to IPv6 on the outgoing interface
- [Multicast policy](#) - used to manage traffic sent to multiple destinations
- [IPv4 access control list](#) - used to filter out packets based on specific IPV4 parameters.
- [IPv6 access control list](#) - used to filter out packets based on specific IPV6 parameters.
- [IPv4 DoS policy](#) - used to prevent malicious or flawed packets on an IPv4 interface from denying access to users.
- [IPv6 DoS policy](#) - used to prevent malicious or flawed packets on an IPv6 interface from denying access to users.

Because the policy determines whether or not NAT will be used, it is also import to look at how to configure:

- [Central SNAT](#) - used for granular controlling when NATing is in use.

Viewing firewall policies

To find a Policy window, follow one of these path in the GUI:

- **Policy & Objects> IPv4 Policy**
- **Policy & Objects> IPv6 Policy**
- **Policy & Objects> NAT64 Policy**
- **Policy & Objects> NAT46 Policy**
- **Policy & Objects> Proxy Policy**
- **Policy & Objects> Multicast Policy**

You may notice other policy options on the left window pane such as:

- **Policy & Objects> IPv4 DoS Policy**
- **Policy & Objects> IPv6 DoS Policy**
- **Policy & Objects> Local InPolicy**

These are different enough that they have their own descriptions in the sections that relate to them.

Menu items

There are some variations, but there are some common elements share by all of them. There is a menu bar across the top. The menu bar will have the following items going from left to right:

- **Create New** button
- **Edit** button
- **Delete** button
- **Search** field
- **Interface Pair View**- Displays the policies in the order that they are checked for matching traffic, grouped by the pairs of Incoming and Outgoing interfaces. For instance, all of the policies referencing traffic from WAN1 to DMZ will be in one section. The policies referencing traffic from DMZ to WAN1 will be in another section. The sections are collapsible so that you only need to look at the sections with policies you are interested in.
- **By Sequence**- Displays the policies in the order that they are checked for matching traffic without any grouping.

Menu items not shared by all policies

- **Policy Lookup** - (IPv4, IPv6)
- **NAT64 Forwarding** - (NAT64)

The Table of Policies

Columns

The tables that make up the Policy window are based on rows which represent individual policies and the columns that represent the various parameters or status within the policy. The columns are customizable by which columns are included and what order they are in.

The table can be laid out a number ways to suit the viewer. There is a column for most of the important pieces of information that you might be interested in seeing, but a lot of them are hidden by default. If you had a large enough screen, you might be able to show all of the columns, but even then it might look a bit busy and cramped together. Figure out which pieces of information are most important to you and hide the rest.

To configure which columns are visible and which are hidden, right click on the header row of the table. This will present a drop down menu. The drop down will be divided into sections. At the top will be the **Selected Columns** which are currently visible, and the next section will be **Available Columns** which show which columns are available to add to the table.

To move a column from the **Available** list to the **Selected** list just click on it. To move a column from the **Selected** list to the **Available** list, it also just takes a click of the mouse. To make the changes show up on the table, go to the bottom of the drop down menu and select **Apply**. Any additions to the table will show up on the right side.

One of the more useful ones that can be added is the ID column. The reason for adding this one is that within the configuration file and CLI, the policies are referenced by their ID number. Some policy settings are only available for configuration in the CLI. If you are looking in the CLI you will see that the only designation for a policy is its number and if you wish to edit the policy or change its order in the sequence you will be asked to move it before or after another policy by referencing its number.

How “Any” policy can remove the Interface Pair View

The FortiGate unit will automatically change the view on the policy list page to **By Sequence** whenever there is a policy containing “any” as the Source or Destination interface. If the **Interface Pair View** is grayed out it is likely that one or more of the policies has used the “any” interface.

By using the “any” interface, the policy should go into multiple sections because it could effectively be any of a number of interface pairings. As mentioned, policies are sectioned by using the interface pairings (for example,

port1 -> port2) and each section has its own specific policy order. The order in which a policy is checked for matching criteria to a packet's information is based solely on the position of the policy within its section or within the entire list of policies as a whole but if the policy is in multiple sections at the same time there is no mechanism for placing the policy in a proper order within all of those sections at the same time because it is a manual process and there is no parameter to compare the precedence of one section or policy over the other. Thus a conflict is created. In order to resolve the conflict the FortiGate firewall removes that aspect of the sections so that there is no need to compare and find precedence between the sections and it therefore has only the Global View to work with.

Policy names

Each policy has a name field. Every policy name must be unique for the current VDOM regardless of policy type. Previous to FortiOS 5.4, this field was optional.



On upgrading from an earlier version of FortiOS to 5.4, policy names are not assigned to old policies, but when configuring new policies, a unique name must be assigned to the policy.

Configuring the Name field

GUI

In the GUI, the field for the policy name is the first field on the editing page.

CLI

In the CLI, the syntax for assigning the policy name is:

```
config firewall [policy|policy6]
  edit 0
    set name <policy name>
  end
```

Disabling policy name requirement

While by default the requirement of having a unique name for each policy is the default, it can be enabled or disabled. Oddly enough, if disabling the requirement is a one time thing, doing it in the CLI is more straightforward.



This setting is VDOM based so if you are running multiple VDOMs, you will have to enter the correct VDOM before entering the CLI commands or turning the feature on or off in the GUI.

GUI

To edit the requirement in the GUI, the ability to do so must be enabled in the CLI. The syntax is:

```
config system settings
  set gui-allow-unnamed-policy [enable|disable]
end
```


Once it has been enabled, the requirement for named policies can be relaxed by going to **System > Feature Visibility**. Allow **Unnamed Policies** can be found under **Additional Features**. Here you can toggle the requirement on and off.

CLI

To change the requirement in the CLI, use the following syntax:

```
config system settings
    set gui-advance policy [enable|disable]
end
```

IPv4 policy

To configure a IPv4 policy in the GUI

1. Go to **Policy & Objects > IPv4 Policy**

The right side window will display a table of the existing IPv4 Policies.

- To edit an existing policy, double click on the policy you wish to edit
- To create a new policy, select the **Create New** icon in the top left side of the right window.

2. Make sure the policy has a name in the **Name** field



By default, a policy is required to have a name, but it is possible to toggle this requirement on or off in the CLI, or in the GUI if you have first enabled the GUI option in the CLI.

- #### 3. Set the **Incoming Interface** parameter by selecting the field with the "+" next to the field label. Selecting the field will slide out a window from the right where you can select from the available interfaces. You can select one or more specific interfaces or you can select the **any** option. Choosing the any option will remove any other interfaces. For more information on interfaces, check the Concepts section called [Interfaces and Zones](#).
- #### 4. Set the **Outgoing Interface** parameter by selecting the field with the "+" next to the field label. (Same rules apply as with the above step.)



Alias names for interfaces, if used, appear in the headings for the Interface Pair View or what used to be called the Section View.

Multiple interfaces or ANY interface can be added to a firewall policy. This feature can be enabled or disabled in the GUI by going to the **System > Feature Select** page and toggling **Multiple Interface Policies**.

When selecting the Incoming or Outgoing interface of a policy, there are a few choices:

- The **ANY** interface (choosing this will remove all other interfaces)
- 1 A single specific interface
- 1 multiple specific interfaces (can be added at the same time or one at a time)



The GUI is intuitive and straightforward on how to do this. Click on the "+" symbol in the interface field and then select the desired interfaces from the side menu. There are a couple of ways to do it in the CLI:

1. Set the interfaces all at once:

```
config firewall policy
edit 0
set srcintf wan1 wan2
end
```

2. Set the first interface and append additional ones:

```
config firewall policy
edit 0
set srcintf wan1
append srcintf wan2
end
```

5. Set the **Source** parameter by selecting the field with the "+" next to the field label. The source in this case is either the source address, source user or source device of the initiating traffic. When the field is selected a window will slide out from the right. Tabs indicating **Address**, **User** or **Device** options are there to help categorize the options along with the option to search. In order to be able to select one of these options it needs to be configured as a firewall object before hand. The "+" icon next to the **Search** field is a shortcut for creating a new firewall object based on the tab that is currently selected. For the **Address** and **Device** tabs, single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
6. Set the **Destination Address** parameter by selecting the field with the "+" next to the field label. This field is similar to the **Source** field but address objects are the only available options to select. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option. For more information on addresses, check the Firewall Objects section called [Addresses](#).
7. Set the **Schedule** parameter by using the drop down menu to select a preconfigured schedule. The "+" icon next to the **Search** field is a shortcut for creating a new schedule object. For more information on addresses, check the Firewall Objects section called [Firewall schedules](#)
8. Set the **Service** parameter by selecting the field with the "+" next to the field label. (Same mechanics for selection apply as with the other similar fields in this window.) Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option. For more information on services, check the Firewall Objects section called [Services and TCP ports](#).
9. Set the **Action** parameter. Select one of the following options for the action:
 - **ACCEPT** - lets the traffic through to the next phase of analysis
 - **DENY** - drops the session

- **LEARN** - collects information about the traffic for future analysis
- **IPsec** - for using with IPsec tunnels

Because the choice of **Action** determines the settings and options below this parameter in the window the rest of the step are associated with a specific **Action**.

Settings if the **ACCEPT** action is selected.

Firewall / Network Options

10. Set the **NAT** parameter by toggling the slider button. (gray means it is disabled)

The NAT setting section is affected by whether or not Central NAT is enabled.

If Central NAT is enabled, the only option in Firewall / Network options will be whether to enable or disable NAT. The rest of the NAT parameters will be set in the Central SNAT page.

If Central NAT is disabled, there are two additional settings in the Policy configuration page.

11. Set the **Fixed Port** parameter by toggling the slider button. (gray means it is disabled)
12. Set the **IP Pool Configuration** by selection one of the options of:
- **Use Outgoing Interface Address**
 - **Use Dynamic IP Pool**

If the **Use Dynamic IP Pool** option is selected, an additional field will appear with the **+** icon. Selecting this field will slide out a window from the right where a preexisting IP Pool can be chosen. One or more IP Pools can be chosen and the **+** icon next to the **Search** field is a shortcut for creating a new IP Pool.

Security Profiles

13. Enabling the **Use Security Profile Group** option will allow the selection of a profile group instead of selecting the individual profiles for the policy.
14. Disable or enable the various **Security Profiles**. Once a Profile has been toggled into the enabled mode a drop down menu will appear for the purpose of choosing a specific profile. Only one profile can be chosen for each profile type. The **+** icon next to the **Search** field in the drop down menu is a shortcut for creating a new profile.

The list of **Security Profiles** available to set includes:

- **AntiVirus**
- **Web Filter**
- **DNS Filter**
- **Application Control**
- **CASI**
- **IPS**
- **Anti-Spam**
- **DLP Sensor**
- **VoIP**
- **ICAP**
- **Web Application Firewall**
- **Proxy Options**
- **SSL/SSH Inspection**

Logging Options

15. Set the **Log Allowed Traffic** parameter by toggling the slider button (gray means it is disabled).

If the **Log Allowed Traffic** setting is enabled, choose whether to log just **Security Events** or **All Sessions** and determine whether or not to keep a record of the packets by toggling the **Capture Packets** setting on or off.

16. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
17. Toggle whether or not to **Enable this policy**. The default is enabled.
18. Select the **OK** button to save the policy.

Settings if the DENY action is selected

Enable the Log Violation Traffic setting by toggling the slider button.

10. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
11. Toggle whether or not to **Enable this policy**. The default is enabled.
12. Select the **OK** button to save the policy.

Settings if the LEARN action is selected

To get more information on the **LEARN** option, read the Learning mode for Firewall policies topic in [What's new for Firewall in 6.0](#)

Firewall / Network Options

10. Set the **NAT** parameter by toggling the slider button. (gray means it is disabled). Unlike the **ACCEPT** option, whether or not Central NAT is enabled or disabled does not affect this settings options.
11. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
12. Toggle whether or not to **Enable this policy**. The default is enabled.
13. Select the **OK** button to save the policy.

Settings if the IPsec action is selected

VPN Tunnel

10. For the VPN Tunnel field, use the drop down menu to select the VPN tunnel that you want the policy associated with.
 11. Toggle the sliding button to enable or disable the option to **Allow traffic to be initiated from the remote site**
- #### Security Profiles

12. Disable or enable the various **Security Profiles**. Once a Profile has been toggled into the enabled mode a drop down menu will appear for the purpose of choosing a specific profile. Only one profile can be chosen for each profile type. The "+" icon next to the **Search** field in the drop down menu is a shortcut for creating a new profile.

The list of **Security Profiles** available to set includes:

- **AntiVirus**
- **Web Filter**
- **DNS Filter**
- **Application Control**
- **CASI**
- **IPS**
- **Anti-Spam**
- **DLP Sensor**

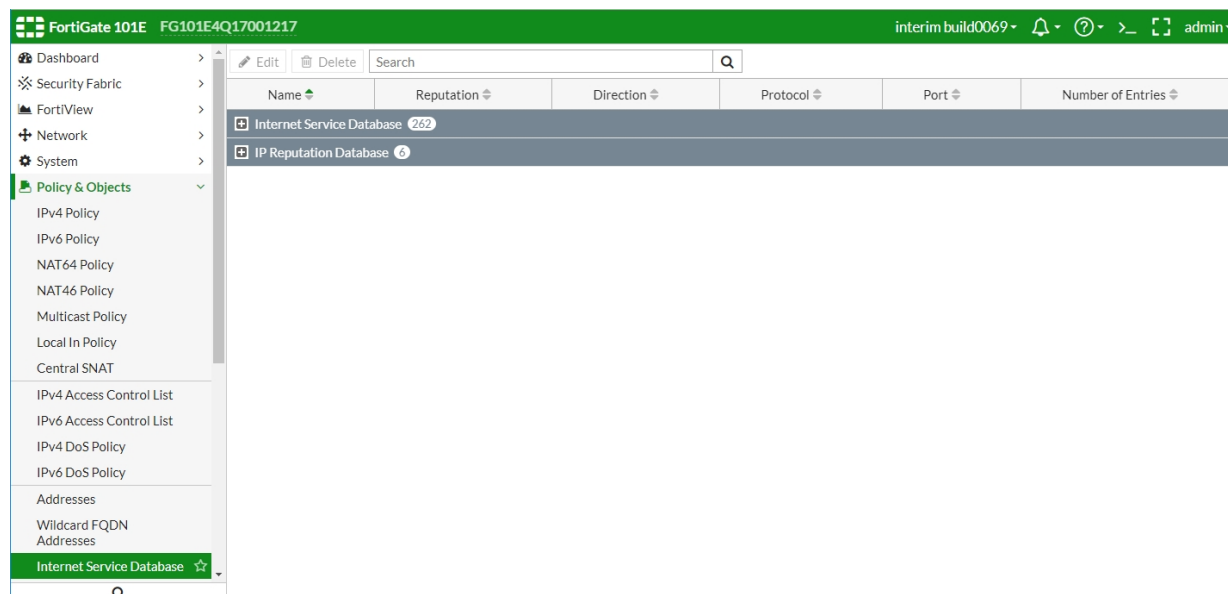
- VoIP
- ICAP
- Web Application Firewall
- Proxy Options
- SSL/SSH Inspection

Logging Options

- Set the **Log Allowed Traffic** parameter by toggling the slider button (gray means it is disabled).
If the **Log Allowed Traffic** setting is enabled, choose whether to log just **Security Events** or **All Sessions** and determine whether or not to keep a record of the packets by toggling the **Capture Packets** setting on or off.
- Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
- Toggle whether or not to **Enable this policy**. The default is enabled.
- Select the **OK** button to save the policy.

ISDB and IRDB in firewall policies

The Internet Service Database (ISDB) and the IP Reputation Database (IRDB) provide similar functionality, so for ease of use, appear together in the GUI.



Use the contents of both, or either database as criteria for inclusion or exclusion in a firewall policy.

Use CLI to define the objects of the ISDB or IRDB objects as parameters within a policy is done within the CLI.

CLI Syntax

```
config firewall policy
edit <ID #>
set internet-service-src {enable|disable}
set internet-service-src-id <ID #>
```

```

set internet-service-src-custom <name>
set internet-service-src-negate {enable|disable}
end

```

CLI options

Option	Description
<code>internet-service-src</code>	Enables or disables the use of Internet Services source for this policy. If enabled, destination address and service are not used.
<code>internet-service-src-id</code>	Internet Service ID Examples: <ul style="list-style-type: none"> • 65536 Google-Others • 65537 Google-Web
<code>internet-service-src-custom</code>	Custom Internet Service name This custom name must already be configured.
<code>internet-service-src-negate</code>	Enables or disables the use of Internet Services in source for this policy. If enabled, <code>internet-service-src</code> specifies what the service must NOT be.



Similar settings are also used in Traffic Shaping policies.

IPv6 policy

To configure a IPv6 policy in the GUI

1. Go to **Policy & Objects > IPv6 Policy**

The right side window will display a table of the existing IPv6 Policies.

- To edit an existing policy, double click on the policy you wish to edit
- To create a new policy, select the **Create New** icon in the top left side of the right window.

2. Make sure the policy has a name in the **Name** field



By default, a policy is required to have a name, but it is possible to toggle this requirement on or off in the CLI, or in the GUI if you have first enabled the GUI option in the CLI.

- #### 3. Set the **Incoming Interface** parameter by selecting the field with the "+" next to the field label. Selecting the field will slide out a window from the right where you can select from the available interfaces. You can select one or more specific interfaces or you can select the **any** option. Choosing the any option will remove any other

interfaces. For more information on interfaces, check the Concepts section called [Interfaces and Zones](#).

4. Set the **Outgoing Interface** parameter by selecting the field with the "+" next to the field label. (Same rules apply as with the above step.)
5. Set the **Source** parameter by selecting the field with the "+" next to the field label. The source in this case is either the source address, source user or source device of the initiating traffic. When the field is selected a window will slide out from the right. Tabs indicating **Address**, **User** or **Device** options are there to help categorize the options along with the option to search. In order to be able to select one of these options it needs to be configured as a firewall object before hand. The "+" icon next to the **Search** field is a shortcut for creating a new firewall object based on the tab that is currently selected. For the **Address** and **Device** tabs, single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
6. Set the **Destination Address** parameter by selecting the field with the "+" next to the field label. This field is similar to the **Source** field but address objects are the only available options to select. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option. For more information on addresses, check the Firewall Objects section called [Addresses](#).
7. Set the **Schedule** parameter by using the drop down menu to select a preconfigured schedule. The "+" icon next to the **Search** field is a shortcut for creating a new schedule object. For more information on addresses, check the Firewall Objects section called [Firewall schedules](#)
8. Set the **Service** parameter by selecting the field with the "+" next to the field label. (Same mechanics for selection apply as with the other similar fields in this window.) Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option. For more information on services, check the Firewall Objects section called [Services and TCP ports](#).
9. Set the **Action** parameter. Select one of the following options for the action:
 - **ACCEPT** - lets the traffic through to the next phase of analysis
 - **DENY** - drops the session

While there are not as many Action options as with the IPv4 policy, because the choice of **Action** determines the settings and options below this parameter in the window the rest of the step are associated with a specific **Action**.

Settings if the **ACCEPT** action is selected.

Firewall / Network Options

10. Set the NAT parameter by toggling the slider button.(gray means it is disabled)
11. Set the **IP Pool Configuration** by selection one of the options of:
 - **Use Outgoing Interface Address**
 - **Use Dynamic IP Pool**

If the **Use Dynamic IP Pool** option is selected:

- An additional field will appear with the + icon. Selecting this field will slide out a window from the right where a preexisting IP Pool can be chosen. One or more IP Pools can be chosen and the "+" icon next to the **Search** field is a shortcut for creating a new IP Pool.
- An additional option to **Preserve the Source Port** will appear as a toggle option. If the slider button is grayed out it is disabled.

Security Profiles

12. Enabling the **Use Security Profile Group** option will allow the selection of a profile group instead of selecting the individual profiles for the policy.

13. Disable or enable the various **Security Profiles**. Once a Profile has been toggled into the enabled mode a drop down menu will appear for the purpose of choosing a specific profile. Only one profile can be chosen for each profile type. The "+" icon next to the **Search** field in the drop down menu is a shortcut for creating a new profile.

The list of **Security Profiles** available to set includes:

- **AntiVirus**
- **Web Filter**
- **Application Control**
- **IPS**
- **Anti-Spam**
- **DLP Sensor**
- **VoIP**
- **ICAP**

Logging Options

14. Set the **Log Allowed Traffic** parameter by toggling the slider button (gray means it is disabled).
If the **Log Allowed Traffic** setting is enabled, choose whether to log just **Security Events** or **All Sessions** and determine whether or not to keep a record of the packets by toggling the **Capture Packets** setting on or off.
15. Add a comment to give a detailed description of the policy in the **Comments** field (up to 1023 characters).
16. Toggle whether or not to **Enable this policy**. The default is enabled.
17. Select the **OK** button to save the policy.

Settings if the DENY action is selected

Enable the **Log Violation Traffic** setting by toggling the slider button.

10. Add a comment to give a detailed description of the policy in the **Comments** field (up to 1023 characters).
11. Toggle whether or not to **Enable this policy**. The default is enabled.
12. Select the **OK** button to save the policy.

NAT64 policy

To configure a NAT64 policy in the GUI

1. Go to **Policy & Objects > NAT64 Policy**

The right side window will display a table of the existing NAT64 Policies.

- To edit an existing policy, double click on the policy you wish to edit
 - To create a new policy, select the **Create New** icon in the top left side of the right window.
2. Set the **Incoming Interface** parameter by selecting the field with the "+" next to the field label. Selecting the field will slide out a window from the right where you can select from the available interfaces. You can select one or more specific interfaces or you can select the **any** option. Choosing the any option will remove any other interfaces. For more information on interfaces, check the Concepts section called [Interfaces and Zones](#).
 3. Set the **Outgoing Interface** parameter by selecting the field with the "+" next to the field label. (Same rules apply as with the above step.)

4. Set the **Source Address** parameter by selecting the field with the "+" next to the field label. The source in this case is an IPv6 Address object of the initiating traffic. When the field is selected a window will slide out from the right. In order to be able to select one of these options it needs to be configured as a firewall object before hand. The "+" icon next to the **Search** field is a shortcut for creating a new firewall object based on the tab that is currently selected. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
5. Set the **Destination Address** parameter by selecting the field with the "+" next to the field label. This field is similar to the **Source Address** field. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option. For more information on addresses, check the Firewall Objects section called [Addresses](#).
6. Set the **Schedule** parameter by using the drop down menu to select a preconfigured schedule. The "+" icon next to the **Search** field is a shortcut for creating a new schedule object. For more information on addresses, check the Firewall Objects section called [Firewall schedules](#)
7. Set the **Service** parameter by selecting the field with the "+" next to the field label. (Same mechanics for selection apply as with the other similar fields in this window.) Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option. For more information on services, check the Firewall Objects section called [Services and TCP ports](#).
8. Set the **Action** parameter. Select one of the following options for the action:
 - **ACCEPT** - lets the traffic through to the next phase of analysis
 - **DENY** - drops the session

While there are not as many Action options as with the IPv4 policy, because the choice of **Action** determines the settings and options below this parameter in the window the rest of the step are associated with a specific **Action**.

Settings if the **ACCEPT** action is selected.

Firewall / Network Options

10. Skip the NAT setting. This type of policy is intended only for traffic that is being NATed from IPv6 to IPv4, because without NATing the traffic couldn't reach its destination, so disabling NAT would be pointless.
11. Set the **IP Pool Configuration** by selection one of the options of:
 - **Use Outgoing Interface Address**
 - **Use Dynamic IP Pool**

If the **Use Dynamic IP Pool** option is selected, an additional field will appear with the + icon. Selecting this field will slide out a window from the right where a preexisting IP Pool can be chosen. One or more IP Pools can be chosen and the "+" icon next to the **Search** field is a shortcut for creating a new IP Pool.

12. Set the **Log Allowed Traffic** parameter by toggling the slider button (gray means it is disabled).
If the **Log Allowed Traffic** setting is enabled, choose whether to log just **Security Events** or **All Sessions** and determine whether or not to keep a record of the packets by toggling the **Capture Packets** setting on or off.
13. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
14. Toggle whether or not to **Enable this policy**. The default is enabled.
15. Select the **OK** button to save the policy.

Settings if the DENY action is selected

Enable the **Log Violation Traffic** setting by toggling the slider button.

10. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
11. Toggle whether or not to **Enable this policy**. The default is enabled.
12. Select the **OK** button to save the policy.

NAT46 policy

To configure a NAT46 policy in the GUI

1. Go to **Policy & Objects > NAT46 Policy**

The right side window will display a table of the existing NAT46 Policies.

- To edit an existing policy, double click on the policy you wish to edit
 - To create a new policy, select the **Create New** icon in the top left side of the right window.
2. Set the **Incoming Interface** parameter by selecting the field with the "+" next to the field label. Selecting the field will slide out a window from the right where you can select from the available interfaces. You can select one or more specific interfaces or you can select the **any** option. Choosing the any option will remove any other interfaces. For more information on interfaces, check the Concepts section called [Interfaces and Zones](#).
 3. Set the **Outgoing Interface** parameter by selecting the field with the "+" next to the field label. (Same rules apply as with the above step.)
 4. Set the **Source** parameter by selecting the field with the "+" next to the field label. The source in this case is either the source address, source user or source device of the initiating traffic. When the field is selected a window will slide out from the right. Tabs indicating **Address**, **User** or **Device** options are there to help categorize the options along with the option to search. In order to be able to select one of these options it needs to be configured as a firewall object before hand. The "+" icon next to the **Search** field is a shortcut for creating a new firewall object based on the tab that is currently selected. For the **Address** and **Device** tabs, single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
 5. Set the **Destination Address** parameter by selecting the field with the "+" next to the field label. This field is similar to the **Source** field but address objects are the only available options to select. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option. For more information on addresses, check the Firewall Objects section called [Addresses](#).
 6. Set the **Schedule** parameter by using the drop down menu to select a preconfigured schedule. The "+" icon next to the **Search** field is a shortcut for creating a new schedule object. For more information on addresses, check the Firewall Objects section called [Firewall schedules](#)
 7. Set the **Service** parameter by selecting the field with the "+" next to the field label. (Same mechanics for selection apply as with the other similar fields in this window.) Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option. For more information on services, check the Firewall Objects section called [Services and TCP ports](#).
 8. Set the **Action** parameter. Select one of the following options for the action:
 - **ACCEPT** - lets the traffic through to the next phase of analysis
 - **DENY** - drops the session

While there are not as many Action options as with the IPv4 policy, because the choice of **Action** determines the settings and options below this parameter in the window the rest of the step are associated with a specific **Action**.

Settings if the **ACCEPT** action is selected.

Firewall / Network Options

10. Skip the NAT setting. This type of policy is intended only for traffic that is being NATed from IPv4 to IPv6, because without NATing the traffic couldn't reach its destination, so disabling NAT would be pointless.
11. Set the **Fixed Port** parameter by toggling the slider button. (gray means it is disabled)
12. Set the **IP Pool Configuration** by selection one of the options of:
 - **Use Outgoing Interface Address**
 - **Use Dynamic IP Pool**

If the **Use Dynamic IP Pool** option is selected, an additional field will appear with the **+** icon. Selecting this field will slide out a window from the right where a preexisting IP Pool can be chosen. One or more IP Pools can be chosen and the **+** icon next to the **Search** field is a shortcut for creating a new IP Pool.

14. Set the **Log Allowed Traffic** parameter by toggling the slider button (gray means it is disabled).
If the **Log Allowed Traffic** setting is enabled, choose whether to log just **Security Events** or **All Sessions** and determine whether or not to keep a record of the packets by toggling the **Capture Packets** setting on or off.
15. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
16. Toggle whether or not to **Enable this policy**. The default is enabled.
17. Select the **OK** button to save the policy.

Settings if the **DENY** action is selected

Enable the **Log Violation Traffic** setting by toggling the slider button.

10. Add a comment to give a detailed description of the policy in the Comments field (up to 1023 characters).
11. Toggle whether or not to **Enable this policy**. The default is enabled.
12. Select the **OK** button to save the policy.

Central SNAT

The Central NAT feature is not enabled by default. When `central-nat` is enabled, `nat` option under IPv4 policies is skipped and SNAT must be done via `central-snat-map`.

- Info messages and redirection links have been added to IPv4 policy list and dialog to indicate the above
- If NGFW mode is policy-based, then it is assumed that central-nat (specifically SNAT) is enabled implicitly
- The option to toggle NAT in central-snat-map policies has been added (previously it was only shown in NGFW policy-based mode).
- In central-snat policy dialog, the port-mapping fields for the original port have been updated to accept ranges.
- Nat will be skipped in firewall policy if per vdom central nat is enabled.
- The Central SNAT window contains a table of all of the Central SNAT policies.

To toggle the feature on or off, use the following commands:

```
config system settings
  set central-nat [enable | disable]
end
```

When Central NAT is enabled the **Central SNAT** section will appear under the Policy & Objects heading in the GUI.

To configure a Central SNAT entry in the GUI

1. Go to **Policy & Objects > Central SNAT**

The right side window will display a table of the existing Central SNAT entries.

- To edit an existing entry, double click on the policy you wish to edit
 - To create a new entry, select the **Create New** icon in the top left side of the right window.
2. Set the **Incoming Interface(s)** by clicking on the "+" in the field. This will slide out a window from the right. Here, you can select from the available interfaces. Selecting a listed interface will highlight it in the window and add it to the field. Clicking on an object in this window while it's highlighted will remove it from the field. Multiple selections are allowed.
 3. Set the **Outgoing Interface(s)** by clicking on the "+" in the field. This will slide out a window from the right. Here, you can select from the available interfaces. Selecting a listed interface will highlight it in the window and add it to the field. Clicking on an object in this window while it's highlighted will remove it from the field. Multiple selections are allowed.
 4. Set the **Source Address** by clicking on the "+" in the field. This will slide out a window from the right. Here, you can select from the available address objects. Selecting a listed object will highlight it in the window and add it to the field. Clicking on an object in this window while it's highlighted will remove it from the field. Multiple selections are allowed. For more information on addresses, check the Firewall Objects section called [Addresses](#).
 5. Set the **Destination Address** by clicking on the "+" in the field. This will slide out a window from the right. Here, you can select from the available address objects. Selecting a listed object will highlight it in the window and add it to the field. Clicking on an object in this window while it's highlighted will remove it from the field. Multiple selections are allowed.

Under the NAT Heading

6. Set the **IP Pool Configuration** parameter by selecting either **Use Outgoing Interface Address** or **Use Dynamic IP Pool**.
 - If Use Dynamic IP Pool is chosen, a field will appear just beneath the option that is used to select which IP Pool object will be used. Set the IP Pool by clicking on the "+" in the field. This will slide out a window from the right. Here, you can select from the available objects.
7. Set the **Protocol** parameter.

There are 5 options for the **Protocol**.

- **ANY** - any protocol traffic
 - **TCP** - TCP traffic only. Protocol number set to 6
 - **UDP** - UDP traffic only. Protocol number set to 17
 - **SCTP** - SCTP traffic only. Protocol number set to 132
 - **Specify** - User can specify the traffic filter protocol by setting the protocol number in the field.
6. If the IP Pool is of the type: Overload, **Explicit Port Mapping** can be enabled.

To enable or disable, use the check box. Once enabled, the following additional parameters will appear.

 - **Original Source Port** - in the left number field, set the starting number of the source port range.
 - **Translated Port** - in the left number field, set the starting number of the translated port range. If it is a single port range leave the right number field alone. If the right number field is set to a number higher than the left, the right number field for the Original Source Port will change to make sure the 2 number ranges have a matching number of ports.
 7. Select the **OK** button to save the entry.

To configure Central SNAT in the CLI

1. Using the CLI interface of your choice, run the following command to get to the correct context.

```
config firewall central-snat-map
```

- To edit an existing entry, run the command `show` or `show full-configuration` to get a listing of all of the entries in the map. Take note of the policy ID for the entry to be edited.
- To create a new entry the next step will use the policy ID 0 which will check for an unused ID number and create an entry with that number.

2. Edit or create an entry with the correct policy ID

```
edit <policyID number>
```

Run the following commands to set the parameters of the entry:

```
set status [enable|disable]
set orig-addr <valid address object preconfigured on the FortiGate>
set srcintf <name of interface on the FortiGate>
set dst-addr <valid address object preconfigured on the FortiGate>
set dstintf <name of interface on the FortiGate>
set protocol <integer for protocol number>
set orig-port <integer for original port number>
set nat-port <integer for translated port number>
set comments <string>
```

3. Save the entry by running the command `end` or `next`.

Example scenarios to showing how CLI treats central-nat

Make nat available regardless of NGFW mode.

```
config firewall central-snat-map
edit 1
set orig-addr "192-86-1-86"
set srcintf "port23"
set dst-addr "192-96-1-96"
set dstintf "port22"
set nat-ippool "pool1"
set protocol 17
set orig-port 2896-2897
set nat enable
end
```

Hide nat-port if nat-ippool is not set or NAT is disabled.

```
config firewall central-snat-map
edit 1
set orig-addr "192-86-1-86"
set srcintf "port23"
set dst-addr "192-96-1-96"
set dstintf "port22"
set nat-ippool "pool1"
set protocol 17
set orig-port 2896-2897
set nat disable
end
```

Change orig-port to accept range

```
config firewall central-snat-map
edit 1
    set orig-addr "192-86-1-86"
    set srcintf "port23"
    set dst-addr "192-96-1-96"
    set dstintf "port22"
    set nat-ippool "pool1"
    set protocol 17
    set orig-port 2896-2897 (help text changed to: Original port or port range).
    set nat-port 35804-35805
end
```

IPv4 access control list

The **IPv4 Access Control List** is a specialized policy for denying IPv4 traffic based on:

- the incoming interface
- the source addresses of the traffic
- the destination addresses of the traffic
- the services or ports the traffic is using

The only action available in this policy is **DENY**

For more information on see [Access Control Lists](#)

To configure a IPv4 access control list entry in the GUI

1. Go to **Policy & Objects > IPv4 Access Control List**

The right side window will display a table of the existing IPv4 Access Control List entries.

- To edit an existing entry, double click on the policy you wish to edit
 - To create a new entry, select the **Create New** icon in the top left side of the right window.
2. Set the **Incoming Interface** parameter by using the drop down menu to select a single interface.
 3. Set the **Source Address** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option. For more information on addresses, check the Firewall Objects section called [Addresses](#).
 4. Set the **Destination Address** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
 5. Set the **Services** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option. For more information on services, check the Firewall Objects section called [Services and TCP ports](#).
 6. Toggle whether or not to **Enable this policy**. The default is enabled.
 7. Select the **OK** button to save the policy.

To configure a IPv4 access control list entry in the CLI

Use the following syntax:

```
config firewall acl
```

```

edit <acl Policy ID #>
    set status enable
    set interface <interface>
    set srcaddr <address object>
    set dstaddr <address object>
    set service <service object>
end
end

```

IPv6 access control list

The **IPv6 Access Control List** is a specialized policy for denying IPv6 traffic based on:

- the incoming interface
- the source addresses of the traffic
- the destination addresses of the traffic
- the services or ports the traffic is using

The only action available in this policy is **DENY**

To configure a IPv6 access control list entry in the GUI

1. Go to **Policy & Objects > IPv6 Access Control List**

The right side window will display a table of the existing IPv6 Access Control List entries.

- To edit an existing entry, double click on the policy you wish to edit
 - To create a new entry, select the **Create New** icon in the top left side of the right window.
2. Set the **Incoming Interface** parameter by using the drop down menu to select a single interface.
 3. Set the **Source IPv6 Address** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option. For more information on addresses, check the Firewall Objects section called [Addresses](#).
 4. Set the **Destination IPv6 Address** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
 5. Set the **Services** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option. For more information on services, check the Firewall Objects section called [Services and TCP ports](#).
 6. Toggle whether or not to **Enable this policy**. The default is enabled.
 7. Select the **OK** button to save the policy.

To configure a IPv6 access control list entry in the CLI

Use the following syntax:

```

config firewall acl6
    edit <acl Policy ID #>
        set status enable
        set interface <interface>
        set srcaddr <address object>
        set dstaddr <address object>
        set service <service object>
    end
end

```

IPv4 DoS policy

To configure a IPv4 DoS policy in the GUI

1. Go to **Policy & Objects > IPv4 DoS Policy**

The right side window will display a table of the existing IPv4 DoS Policies.

- To edit an existing policy, double click on the policy you wish to edit
- To create a new policy, select the **Create New** icon in the top left side of the right window.

2. Set the **Incoming Interface** parameter by using the drop down menu to select a single interface.

3. Set the **Source Address** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option. For more information on addresses, check the Firewall Objects section called [Addresses](#).

4. Set the **Destination Address** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.

5. Set the **Services** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option. For more information on services, check the Firewall Objects section called [Services and TCP ports](#).

6. Set the parameters for the various traffic anomalies.

All of the anomalies that profiles have been created for are in 2 tables. These tables break up the anomaly profiles into **L3 Anomalies** and **L4 Anomalies**. All of the anomalies have the following parameters that can be set on a per anomaly or per column basis.

- Status - enable or disable the indicated profile
- Logging - enable or disable logging of the indicated profile being triggered
- Action - whether to Pass or Block traffic when the threshold is reached
- Threshold - the number of anomalous packets detected before triggering the action.

The listing of anomaly profiles includes:

L3 Anomalies

- ip_src_session
- ip_dst_session

L4 Anomalies

- tcp_syn_flood
- tcp_port_scan
- tcp_src_session
- tcp_dst_session
- udp_flood
- udp_scan
- udp_src_session
- udp_dst_session
- icmp_flood
- icmp_sweep
- icmp_src_session

- sctp_flood
- sctp_scan
- sctp_src_session
- sctp_dst_session

7. Toggle whether or not to **Enable this policy**. The default is enabled.
8. Select the **OK** button to save the policy.

Example

The company wishes to protect against Denial of Service attack. They have chosen some where they wish to block the attacks of the incidence goes above a certain threshold and for some others they are just trying to get a baseline of activity for those types of attacks so they are letting the traffic pass through without action.

- The interface to the Internet is on WAN1
- There is no requirement to specify which addresses are being protected or protected from.
- The protection is to extend to all services.
- The TCP attacks are to be blocked
- The UDP, ICMP, and IP attacks are to be recorded but not blocked.
- The SCTP attack filters are disabled
- The tcp_syn_flood attack's threshold is to be changed from the default to 1000

Configuring the DoS policy in the GUI

1. Go to **Policy & Objects > Policy > DoS**.
2. Create a new policy
3. Fill out the fields with the following information:

Field	Value
Incoming Interface	wan1
Source Address	all
Destination Addresses	all
Service	ALL

L3 Anomalies

Name	Status	Logging	Action	Threshold
ip_src_session	enabled	enabled	Pass	5000
ip_dst_session	enabled	enabled	Pass	5000

L4 Anomalies

Name	Status	Logging	Action	Threshold
tcp_syn_flood	enabled	enabled	Block	1000
tcp_port_scan	enabled	enabled	Block	<default value>
tcp_src_session	enabled	enabled	Block	<default value>
tcp_dst_session	enabled	enabled	Block	<default value>
udp_flood	enabled	enabled	Pass	<default value>
udp_scan	enabled	enabled	Pass	<default value>
udp_src_session	enabled	enabled	Pass	<default value>
udp_dst_session	enabled	enabled	Pass	<default value>
icmp_flood	enabled	enabled	Pass	<default value>
icmp_sweep	enabled	enabled	Pass	<default value>
icmp_src_session	enabled	enabled	Pass	<default value>
icmp_dst_session	enabled	enabled	Pass	<default value>
sctp_flood	not enabled	not enabled	Pass	<default value>
sctp_scan	not enabled	not enabled	Pass	<default value>
sctp_src_session	not enabled	not enabled	Pass	<default value>
sctp_dst_session	not enabled	not enabled	Pass	<default value>

4. Toggle the button next to **Enable this policy** to **ON**.

5. Select **OK**.

Configuring the IPv4 DoS policy in the GUI

Using the CLI of your choice, enter the following commands:

```
config firewall DoS-policy
edit 0
set status enable
set interface wan1
set srcaddr all
set dstaddr all
set service ALL
config anomaly
edit "tcp_syn_flood"
set status enable
```

```
        set log disable
        set action block
        set threshold 1000
    next
edit "tcp_port_scan"
    set status enable
    set log disable
    set action block
    set threshold 1000
next
edit "tcp_src_session"
    set status enable
    set log disable
    set action block
    set threshold 5000
next
edit "tcp_dst_session"
    set status enable
    set log disable
    set action block
    set threshold 5000
next
edit "udp_flood"
    set status enable
    set log disable
    set action pass
    set threshold 2000
next
edit "udp_scan"
    set status enable
    set log disable
    set action pass
    set quarantine none
    set threshold 2000
next
edit "udp_src_session"
    set status enable
    set log disable
    set action pass
    set threshold 5000
next
edit "udp_dst_session"
    set status enable
    set log disable
    set action pass
    set threshold 5000
next
edit "icmp_flood"
    set status enable
    set log disable
    set action pass
    set threshold 250
next
edit "icmp_sweep"
    set status enable
    set log disable
    set action pass
```

```
        set threshold 100
      next
    edit "icmp_src_session"
      set status enable
      set log disable
      set action pass
      set threshold 300
    next
    edit "icmp_dst_session"
      set status enable
      set log disable
      set action pass
      set threshold 1000
    next
    edit "ip_src_session"
      set status disable
      set log enable
      set action pass
      set threshold 5000
    next
    edit "ip_dst_session"
      set status disable
      set log enable
      set action pass
      set threshold 5000
    next
    edit "sctp_flood"
      set status disable
      set log disable
      set action pass
      set threshold 2000
    next
    edit "sctp_scan"
      set status disable
      set log disable
      set action pass
      set threshold 1000
    next
    edit "sctp_src_session"
      set status disable
      set log disable
      set action pass
      set threshold 5000
    next
    edit "sctp_dst_session"
      set status disable
      set log disable
      set action pass
      set threshold 5000
    next
  end
end
end
```



In this example of the CLI, all of the relevant settings have been left in, but some of them are default settings and would not have to have been specifically set to work. For instance, if the action parameter is not set it automatically defaults to pass.

IPv6 DoS policy

To configure a IPv6 DoS policy in the GUI

1. Go to **Policy & Objects > IPv6 DoS Policy**

The right side window will display a table of the existing IPv6 DoS Policies.

- To edit an existing policy, double click on the policy you wish to edit
 - To create a new policy, select the **Create New** icon in the top left side of the right window.
2. Set the **Incoming Interface** parameter by using the drop down menu to select a single interface.
 3. Set the **Source IPv6 Address** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option. For more information on addresses, check the Firewall Objects section called [Addresses](#).
 4. Set the **Destination IPv6 Address** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
 5. Set the **Services** parameter by selecting the field with the "+" next to the field label. Single or multiple options can be selected unless the **ALL** option is chosen in which case, it will be the only option. For more information on services, check the Firewall Objects section called [Services and TCP ports](#).
 6. Set the parameters for the various traffic anomalies.

All of the anomalies that profiles have been created for are in 2 tables. These tables break up the anomaly profiles into **L3 Anomalies** and **L4 Anomalies**. All of the anomalies have the following parameters that can be set on a per anomaly or per column basis.

- Status - enable or disable the indicated profile
- Logging - enable or disable logging of the indicated profile being triggered
- Action - whether to Pass or Block traffic when the threshold is reached
- Threshold - the number of anomalous packets detected before triggering the action.

The listing of anomaly profiles includes:

L3 Anomalies

- ip_src_session
- ip_dst_session

L4 Anomalies

- tcp_syn_flood
- tcp_port_scan
- tcp_src_session
- tcp_dst_session
- udp_flood
- udp_scan
- udp_src_session

- udp_dst_session
 - icmp_flood
 - icmp_sweep
 - icmp_src_session
 - icmp_dst_session
 - sctp_flood
 - sctp_scan
7. Toggle whether or not to **Enable this policy**. The default is enabled.
 8. Select the **OK** button to save the policy.

Configuring the IPv6 DoS policy in the GUI

The configuring of the IPv6 version of the DoS policy is the same as in the IPv4 version , with the exception of first command.

Using the CLI of your choice, enter the following commands:

```
config firewall DoS-policy6
```

The rest of the settings are the same as in IPv4 Dos Policy.

Multicast policy

The **Multicast Policy** GUI page has been updated from previous versions of the firmware to the new GUI look and feel. Some functionality has also been changed.

The DNAT option has been removed from the GUI but is still in the CLI.

To create/edit a multicast policy go to **Policy & Objects > Multicast Policy**. The Listing window on the right will have buttons along the top that will enable you to

- **Create New**
- **Edit**
- **Delete**

There is also a **Search** field that will allow you to search or filter the available policies if you have a lot of them.

To configure a new policy left click on the **Create New** button. This will reveal the New Policy editing window.

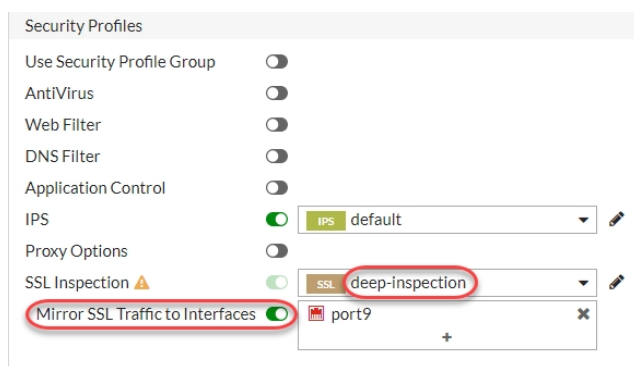
1. Using the drop down menu, fill in the field for **Incoming Interface**. Only one interface can be chosen.
2. Using the drop down menu, fill in the field for **Outgoing Interface**. Only one interface can be chosen.
3. Set the **Source Address** parameter by selecting the field with the "+" next to the field label. When the field is selected a window will slide out from the right. In order for a multicast address to be available for selection, the address object needs to have been created already. Only useable address options will be available for selection. This means only multicast address objects and the more generic **all** and **none** options. The "+" icon next to the **Search** field is a shortcut for creating a new firewall object. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
4. Set the **Destination Address** parameter by selecting the field with the "+" next to the field label. This field is similar to the **Source** field. Single or multiple options can be selected unless the **all** option is chosen in which case, it will be the only option.
5. Set the **Action** parameter. This will be to either **ACCEPT** or **DENY** the traffic through the policy.

6. Toggle the **Enable SNAT** switch to the setting you want. If the slider is gray the option is disabled. If it is colored, it is enabled.
7. Use the drop down menu to select a **Protocol**. The options are:
 - **Any**
 - **ICMP**
 - **IGMP**
 - **TCP** - includes **Port Range** fields
 - **UDP** - includes **Port Range** fields
 - **OSPF**
 - **Other** - includes a field for the protocol number
8. Depending on which Protocol is defined, the some other fields may appear.
 - **Port Range** - The first field is for the starting value for the port and the second for the ending value for the port range used by the protocol. Both of these values are inclusive.
 - Protocol field - This appears when the **Other** option is chosen. Enter the value of the protocol number for the protocol you wish to use.
9. Toggle the **Log Allowed Traffic** switch to the setting you want. If the slider is gray the option is disabled. If it is colored, it is enabled.
10. Toggle the **Enable this policy** switch to the setting you want. If the slider is gray the option is disabled. If it is colored, it is enabled. By default, this should be enabled
11. Click on the **OK** button to save the policy.

SSL mirroring for policies

You can configure the mirroring of SSL inspected traffic for IPv4 and IPv6 policies in the GUI. To use this option, configure the FortiGate to use SSL deep inspection instead of certificate inspection.

SSL inspection is automatically enabled when you enable a security profile in the policy configuration page. Selecting an SSL inspection profile that uses deep inspection reveals the **Mirror SSL Traffic to Interfaces** option. If you enable mirroring, an **SSL Mirror Terms of Use** agreement appears. Once you accept the agreement, you can select interfaces to mirror the traffic to in the Mirror SSL Traffic to Interfaces field.



Object configuration

As was mentioned earlier, the components of the FortiGate firewall go together like interlocking building blocks. The Firewall objects are a prime example of those building blocks. They are something that can be configured once and then used over and over again to build what you need. They can assist in making the administration of the FortiGate unit easier and more intuitive as well as easier to change. By configuring these objects with their future use in mind as well as building in accurate descriptions the firewall will become almost self documenting. That way, months later when a situation changes, you can take a look at a policy that needs to change and use a different firewall object to adapt to the new situation rather than build everything new from the ground up to accommodate the change.

This chapter includes information about the following Firewall objects:

- [Addresses](#)
- ["Virtual IPs" on page 728](#)
- [IP Pools](#)
- ["Services" on page 742](#)
- ["Firewall schedules" on page 749](#)

UUID support

A Universally Unique Identified (UUID) attribute has been added to some firewall objects, so that the logs can record these UUID to be used by a FortiManager or FortiAnalyzer unit. The objects currently include:

- Addresses, both IPv4 and IPv6
- Address Groups, both IPv4 and IPv6
- Virtual IPs, both IPv4 and IPv6
- Virtual IP groups, both IPv4 and IPv6
- Policies, IPv4, IPv6 and IP64

A UUID is a 16-octet (128-bit) number that is represented by 32 lowercase hexadecimal digits. The digits are displayed in five groups separated by hyphens (-). The pattern is 8-4-4-4-12; 36 digits if you include the hyphens.



Note: UUID is only supported on large-partition platforms ($\geq 128M$)

Addresses

Firewall addresses define sources and destinations of network traffic and are used when creating policies. When properly set up these firewall objects can be used with great flexibility to make the configuration of firewall policies simpler and more intuitive. The FortiGate unit compares the IP addresses contained in packet headers with a security policy's source and destination addresses to determine if the security policy matches the traffic.

The address categories and the types within those categories on the FortiGate unit can include:

- IPv4 addresses
 - IP address and Netmask
 - IP address range

- Geography based address
- Fully Qualified Domain Name (FQDN) address
- Wildcard FQDN
- IPv4 address group
- IPv6 addresses
 - Subnets
 - IP range
 - IPv6 address group
- Multicast addresses
 - Multicast IP range
 - Broadcast subnets
- Proxy addresses
 - URL pattern
 - Host Regex match
 - URL category
 - Http method
 - User agent
 - HTTP header
 - Advanced (source)
 - Advanced (destination)
- IP Pools (IPv4)
 - Overload
 - One-to-one
 - Fixed port range
 - Port block allocation
- IP pools (IPv6)
- Virtual IP addresses
 - IPv4
 - IPv6
 - NAT46
 - NAT64

Interfaces

When setting up an address one of the parameters that is asked for is the interface. This means that the system will expect to see that address only on the interface that you select. You can only select one interface. If you expect that the address may be seen at more than one interface you can choose the “any” interface option. Whenever, possible it is best to choose a more specific interface than the “any” option because in the GUI configuration of firewall policies there is a drop down field that will show the possible addresses that can be used. The drop down will only show those addresses that can be on the interface assigned for that interface in the policy.

Example:

- You have an address called “XYZ”.
- “XYZ” is set to the WAN1 interface because that is the only interface that will be able to access that address.

- When you are selecting a Source Address in the Web-based Manager for a policy that is using the DMZ the address “XYZ” will not be in the drop-down menu.

When there are only 10 or 20 addresses this is not a concern, but if there are a few hundred addresses configured it can make your life easier.

Addresses, address groups, and virtual IPs must have unique names to avoid confusion in firewall policies. If an address is selected in a policy, the address cannot be deleted until it is deselected from the policy.

Addressing Best Practices Tip



The other reason to assign a specific interface to addresses is that it will prevent you from accidentally assigning an address where it will not work properly. Using the example from earlier, if the “XYZ” address was assigned to the “Any” interface instead of WAN1 and you configure the “XYZ” address.

Addressing Best Practices Tip



Don't specify an interface for VIP objects or other address objects that may need to be moved or approached from a different direction. When configuring a VIP you may think that it will only be associated with a single interface, but you may later find that you need to reference it on another interface.

Example: Some web applications require the use of a FQDN rather than an IP address. If you have a VIP set up that works from the Internet to the Internal LAN you won't be able to use that VIP object to access it from an internal LAN interface.

IPv4 addresses

When creating an IPv4 address there are a number of different types of addresses that can be specified. These include:

- FQDN
- Geography
- IP range
- IP/Netmask
- Wildcard FQDN

Which one chosen will depend on which method most easily yet accurately describes the addresses that you are trying to include with as few entries as possible based on the information that you have. For instance, if you are trying to describe the addresses of a specific company's web server but if you have no idea of how extensive there web server farm is you would be more likely to use a Fully Qualified Domain Name (FQDN) rather than a specific IP address. On the other hand some computers don't have FQDNs and a specific IP address must be used.

The following is a more comprehensive description of the different types of addresses.

FQDN addresses

By using Fully Qualified Domain Name (FQDN) addressing you can take advantage of the dynamic ability of DNS to keep up with address changes without having to manually change the addresses on the FortiGate. FQDN addresses are most often used with external web sites but they can be used for internal web sites as well if there is a trusted DNS server that can be accessed. FQDN addressing also comes in handy for large web sites that may use multiple addresses and load balancers for their web sites. The FortiGate firewall automatically maintains a cached record of all the addresses resolved by the DNS for the FQDN addresses used.

For example, if you were doing this manually and you wanted to have a security policy that involved Google you could track down all of the IP addresses that they use across multiple countries. Using the FQDN address is simpler and more convenient.

When representing hosts by an FQDN, the domain name can also be a subdomain, such as mail.example.com.

Valid FQDN formats include:

- <host_name>.<top_level_domain_name> such as example.com
- <host_name>.<second_level_domain_name>.<top_level_domain_name>, such as mail.example.com

When creating FQDN entries it is important to remember that:

- Wildcards are not supported in FQDN address objects
- While there is a level of convention that would imply it, “www.example.com” is not necessarily the same address of “example.com”. they will each have their own records on the DNS server.

The FortiGate firewall keeps track of the DNS TTLs so as the entries change on the DNS servers the IP address will effectively be updated for the FortiGate. As long as the FQDN address is used in a security policy, it stores the address in the DNS cache.



There is a possible security downside to using FQDN addresses. Using a fully qualified domain name in a security policy means that your policies are relying on the DNS server to be accurate and correct. DNS servers in the past were not seen as potential targets because the thinking was that there was little of value on them and therefore are often not as well protected as some other network resources. People are becoming more aware that the value of the DNS server is that in many ways it controls where users and computers go on the Internet. Should the DNS server be compromised, security policies requiring domain name resolution may no longer function properly.

Creating a Fully Qualified Domain Name address

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**. A drop down menu is displayed. Select **Address**.
3. In the **Category** field, chose **Address**. (This is for IPv4 addresses.)
4. Input a **Name** for the address object.
5. In the **Type** field, select **FQDN** from the drop down menu.
6. Input the domain name in the **FQDN** field.
7. In the **Interface** field, leave as the default **any** or select a specific interface from the drop down menu.
8. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled, the address will appear in drop down menus where it is an option.
9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example: FQDN address

You have to create a policy that will govern traffic that goes to a site that has a number of servers on the Internet. Depending on the traffic or the possibility that one of the servers is down network traffic can go to any one of those sites. The consistent factor is that they all use the same Fully Qualified Domain Name.

- The FQDN of the web site: example.com
- The number of ISP connections off of the FortiGate firewall: 2

Configuring the address in the GUI

1. Go to **Policy & Objects > Objects > Addresses** and select **Create New > Address**.
2. Fill out the fields with the following information:

Category	Address
Name	BigWebsite.com
Type	FQDN
FQDN	bigwebsite.com
Interface	any
Show in Address List	<enable>
Comments	<Input into this field is optional>

3. Select **OK**.

Configuring the address in the CLI

```
config firewall address
edit BigWebsite.com
set type fqdn
set associated-interface any
set fqdn bigwebsite.com
end
```

Verification

To verify that the addresses were added correctly:

1. Go to **Firewall Objects > Address > Addresses**. Check that the addresses have been added to the address list and that they are correct.
2. Enter the following CLI command:

```
config firewall address
edit <the name of the address that you wish to verify>
Show full-configuration
```

Changing the TTL of a FQDN address

To make sure that the FQDN resolves to the most recent active server you have been asked to make sure that the FortiGate has not cached the address for any longer than 10 minutes.

There is no field for the cached time-to-live in the web-based manager. It is only configurable in the CLI. Enter the following commands:

```
config firewall address
edit BigWebsite.com
set cache-ttl 600
end
```

Geography based addresses

Geography addresses are those determined by country of origin.

This type of address is only available in the IPv4 address category.

Creating a geography address

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**. A drop down menu is displayed. Select **Address**.
3. In the **Category** field, chose **Address**. (This is for IPv4 addresses.)
4. Input a **Name** for the address object.
5. In the **Type** field, select **Geography** from the drop down menu.
6. In the **Country** field, select a single country from the drop down menu.
7. In the **Interface** field, leave as the default **any** or select a specific interface from the drop down menu.
8. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example: Geography-based address

Configuring the address in the GUI

Your company is US based and has information on its web site that may be considered information that is not allowed to be sent to embargoed countries. In an effort to help reduce the possibility of sensitive information going to those countries you have been asked to set up addresses for those countries so that they can be blocked in the firewall policies.

- One of the countries you have been asked to block is Cuba
- You have been asked to comment the addresses so that other administrators will know why they have been created

1. Go to **Policy & Objects > Objects > Addresses** and select **Create New > Address**.
2. Fill out the fields with the following information

Category	Address
Name	Cuba
Type	Geography

Country	Cuba
Interface	any
Visibility	<enable>
Comments	Embargoed

3. Select **OK**.

Configuring the address in the CLI

Enter the following CLI commands:

```
config firewall address
edit Cuba
set type geography
set country CN
set interface wan1
end
```

Overrides

It is possible to assign a specific ip address range to a customized country ID. Generally, geographic addressing is done at the VDOM level; it could be considered global if you are using the root VDOM, but the geoip-override setting is a global setting.

```
config system geoip-override
edit "test"
set country-id "A0"
config ip-range
edit 1
set start-ip 7.7.7.7
set end-ip 7.7.7.8
next
edit 2
set start-ip 7.7.10.1
set end-ip 7.7.10.255
end
```



- While the setting exists in the configuration file, the system assigns the country-id option automatically.
- While you can use "edit 1" and "edit 2", it is simpler to use "edit 0" and let the system automatically assign an ID number.

After creating a customized Country by using geoip-override command, the New country name has been added automatically to the country list and will be available on the Firewall Address Country field.

Diagnose commands

There are a few diagnose commands used with geographic addresses. The basic syntax is:

```
diagnose firewall ipgeo [country-list | ip-list | ip2country | override |
copyright-notice]
```

Diagnose command	Description
country-list	Listing of all the countries.
ip-list	List of the IP addresses associated with the country
ip2country	Used to determine which country a specific IP address is assigned to.
override	Listing of user defined geography data - items configured by using "config system geoip-override" command.
copyright-notice	Shows the copyright notice.



Click on the diagnose command in the table to connect to the Fortinet Diagnose Wiki page that deals with the command option, to get more information.

IP range addresses

Where the subnet address is good at representing a standardized group of addresses that are subnets the IP Range type of address can describe a group of addresses while being specific and granular. It does this by specifying a continuous set of IP addresses between one specific IP address and another. While it is most common that this range is with a subnet it is not a requirement. For instance, 192.168.1.0/24 and 192.168.2.0/24 would be 2 separate subnets but if you wanted to describe the top half of one and the bottom half of the other you could describe the range of 192.168.1.128-192.168.2.127. It's also a lot easier than trying to calculate the correct subnet mask.

The format would be:

x.x.x.x-x.x.x.x, such as 192.168.110.100-192.168.110.120

There is a notation that is commonly used and accepted by some devices that follows the format:

x.x.x.[x-x], such as 192.168.110.[100-120]

This format is not recognized in FortiOS 5.2 as a valid IP Range.

Creating a IP range address

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**. A drop down menu is displayed. Select **Address**
3. In the **Category** field, chose **Address**(IPv4 addresses) or **IPv6 Address**.
4. Input a **Name** for the address object.
5. In the **Type** field, select **IP Range** from the drop down menu.

6. In the **Subnet / IP Range** field, enter the range of addresses in the following format: x.x.x.x-x.x.x.x (no spaces)
7. In the **Interface** field, leave as the default **any** or select a specific interface from the drop down menu. (This setting is not available for IPv6 addresses)
8. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example

Example of a IP Range address for a group of computers set aside for guests on the company network.

Field	Value
Category	Address or IPv6 Address
Name	Guest_users
Type	IP Range
Subnet / IP Range	192.168.100.200-192.168.100.240
Interface	Port1
Show in Address List	[on]
Comments	Computers on the 1st floor used by guests for Internet access.



IP Range addresses can be configured for both IPv4 and IPv6 addresses. The only differences in creating an IPv6 IP Range address is that you would choose IPv6 Address for the Category and the syntax of the address in the Subnet/IP Range field would be in the format of 2001:0db8:0000:0002:0:0:0:20-2001:0db8:0000:0004:0:0:0:20

IP / netmask addresses

The subnet type of address is expressed using a host address and a subnet mask. From a strictly mathematical stand point this is the most flexible of the types because the address can refer to as little one individual address or as many as all of the available addresses.

It is usually used when referring to your own internal addresses because you know what they are and they are usually administered in groups that are nicely differentiated along the lines of the old A, B, and C classes of IPv4 addresses. They are also addresses that are not likely to change with the changing of Internet Service Providers (ISP).

When representing hosts by an IP address with a netmask, the IP address can represent one or more hosts. For example, a firewall address can be:

- A single host such as a single computer with the address 192.45.46.45
- A range of hosts such as all of the hosts on the subnet 192.45.46.1 to 192.45.46.255
- All hosts, represented by 0.0.0.0 which matches any IP address

The netmask corresponds to the subnet class of the address being added, and can be represented in either dotted decimal or CIDR format. The FortiGate unit automatically converts CIDR formatted netmasks to dotted decimal format. Example formats:

- Netmask for a class A subnet of 16,777,214 usable addresses: 255.0.0.0, or /8
- Netmask for a class B subnet of 65,534 usable addresses: 255.255.0.0, or /16
- Netmask for a class C subnet of 254 usable addresses: 255.255.255.0, or /24
- Netmask for subnetted class C of 126 usable addresses: 255.255.255.128, or /25
- Netmask for subnetted class C of 62 usable addresses: 255.255.255.128, or /26
- Netmask for subnetted class C of 30 usable addresses: 255.255.255.128, or /27
- Netmask for subnetted class C of 14 usable addresses: 255.255.255.128, or /28
- Netmask for subnetted class C of 6 usable addresses: 255.255.255.128, or /29
- Netmask for subnetted class C of 2 usable addresses: 255.255.255.128, or /30
- Netmask for a single computer: 255.255.255.255, or /32
- Netmask used with 0.0.0.0 to include all IP addresses: 0.0.0.0, or /0

So for a single host or subnet the valid format of IP address and netmask could be either:

x.x.x.x/x.x.x.x, such as 192.168.1.0/255.255.255.0

or

x.x.x.x/x, such as 192.168.1.0/24

Static route configuration

A setting that is found in the IP/Netmask address type that is not found in the other address types is the enabling or disabling of **Static Route Configuration**. Enabling this feature includes the address in the listing of named addresses when setting up a static route.

To use in the GUI

1. Enable the **Static Route Configuration** in the address.
2. Go to **Network > Static Routes** and create a new route.
3. For a **Destination** type, choose **Named Address**.
4. Using the drop down menu, enter the name of the address object in the field just underneath the **Destination** type options.
5. Fill out the other information relevant to the route
6. Select the **OK** button

To enable in the CLI:

```
config firewall address
  edit <address_name>
    set allow-routing enable
  end
```

Creating a subnet address

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**. A drop down menu is displayed. Select **Address**.
3. In the **Category** field, chose **Address**. (This is for IPv4 addresses.)
4. Input a **Name** for the address object.
5. In the **Type** field, select **IP/Netmask** from the drop down menu.
6. In the **Subnet/IP Range** field, enter the address and subnet mask according to the format x.x.x.x/x.x.x.x or the short hand format of x.x.x.x/x
7. In the **Interface** field, leave as the default **any** or select a specific interface from the drop down menu.
8. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
9. Select the desired on/off toggle setting for **Static Route Configuration**.
10. Input any additional information in the **Comments** field.
11. Press **OK**.

Example

Example of a Subnet address for a database server on the DMZ:

Field	Value
Category	Address
Name	DB_server_1
Type	IP/Netmask
Subnet/IP Range	United States
Interface	any
Show in Address List	[on]
Static Route Configuration	[off]
Comments	

Wildcard FQDN

There are a number of companies that use secondary and even tertiary domain names or FQDNs for their websites. Wildcard FQDN addresses are to ease the administrative overhead in cases where this occurs. Sometimes its as simple as sites that still use www. as a prefix for their domain name. If you don't know whether or not the www is being used it's simpler to use a wildcard and include all of the possibilities whether it be example.com, www.example.com or even ftp.example.com.

The following wildcard character instances are supported in wildcard FQDN addresses:

- "?" character
- "*" character in the middle of a phrase

- The "?*" combination



Wildcard FQDN addresses do not resolve to a specific set of IP addresses in the same way that a normal FQDN address does. They are intended for use in SSL exemptions and should not be used as source or destination addresses in policies.

Creating a Fully Qualified Domain Name address

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**. A drop down menu is displayed. Select **Address**
3. In the **Category** field, chose **Address**. (This is for IPv4 addresses.)
4. Input a **Name** for the address object.
5. In the **Type** field, select **Wildcard FQDN** from the drop down menu.
6. Input the domain name in the **Wildcard FQDN** field.
7. In the **Interface** field, leave as the default **any** or select a specific interface from the drop down menu.
8. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example

Example of a FQDN address for a remote FTP server used by Accounting team:

Field	Value
Category	Address
Name	Example.com_servers
Type	Wildcard FQDN
Wildcard FQDN	*.example.com
Interface	any
Show in Address List	[on]
Comments	Secondary and tertiary domain names for example.com

Wildcard FQDNs for SSL deep inspection exemptions

As part of an improvement to SSL deep inspection, wild card FQDN addresses are stored in two tables, one relates to `firewall address`, historic location for the information, and the second location relates to `firewall wildcard-fqdn custom`. The wildcard FQDN in `firewall address` is used by `proxy-policy`. The wildcard FQDN in `firewall wildcard-fqdn custom` is used by `ssl-exempt` in `ssl-ssh-profile`.



During an upgrade from v5 to v6, all wildcard FQDN in firewall address in the v5 configuration will be moved to firewall wildcard-fqdn custom. If the wildcard FQDN is used in a policy in v5, the upgrade process will leave a copy of the wildcard FQDN in firewall address in addition to the one in firewall wildcard-fqdn custom.

Syntax of the firewall wildcard-fqdn custom object:

```
config firewall wildcard-fqdn custom
  edit <string_value>
    set uuid <string_value>
    set wildcard-fqdn <string_value>
    set color <integer 0-32>
    set comment <string_value>
    set visibility {enable|disable}
  next
end
```

Syntax of the firewall wildcard-fqdn group object:

```
config firewall wildcard-fqdn group
  edit "test-group"
    set uuid <string_value>
    set member <string_value> [<string_value>]
    set color 0
    set comment ''
    set visibility enable
  next
end
```



In the CLI, separate group members with a space.

IPv6 addresses

When creating an IPv6 address there are a number of different types of addresses that can be specified. These include:

- Subnet
- IP Range - the details of this type of address are the same as the IPv4 version of this type
- IPv6 FQDN firewall addresses - similar to the IPv4 version.

The IPv6 addresses don't yet have the versatility of the IPv4 address in that they don't have things like geography based addresses, but as IPv6 becomes more mainstream this should change.

Subnet addresses

The Subnet Address type is one that is only used in reference to IPv6 addresses. It represents an IPv6 address subnet. This means that the address will likely be a series of hexadecimal characters followed by a double colon, followed by a "/", and then a number less than 128 to indicate the size of the subnet. An example would be:

fd5e:3c59:35ce:f67e::/64

- The hexadecimal characters represent the IPv6 subnet address.
- The "::" indicates 0's from that point to the left. In an actual address for a computer, the hexadecimal characters that would take the place of these zeros would represent the device address on the subnet.
- /xx, in this case /64 represents the number of bits in the subnet. This will make a range that can potentially include 18,446,744,073,709,551,616 addresses. For those wanting to use English rather than math, that is 18 Quintillion.

Creating a subnet address

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**. A drop down menu is displayed. Select **Address**
3. In the **Category** field, chose **IPv6 Address**.
4. Input a **Name** for the address object.
5. In the **Type** field, select **Subnet** from the drop down menu.
6. In the **Subnet / IP Range** field, enter the range of addresses in IPv6 format (no spaces)
7. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
8. Input any additional information in the **Comments** field.
9. Press **OK**.

Example

Example of a IP Range address for a group of computers set aside for guests on the company network.

Field	Value
Category	IPv6 Address
Name	IPv6_Guest_user_range
Type	Subnet
Subnet / IP Range	fd5e:3c59:35ce:f67e::/64
Show in Address List	[on]
Comments	

IPv6 FQDN firewall addresses

FQDN firewall addresses can be configured for IPv6.

Syntax in CLI

```
config firewall address6
  edit <address_name>
    set type fqdn
    set fqdn <domain_name>
    set cache-ttl <integer value from 0 to 86400>
  end
```

Firewall IPv6 address templates

You can use the IPv6 address templates to create new IPv6 addresses that share a prefix. Using templates for addresses reduces the chance of configuring an incorrect address due to a typographical error.

- A standard IPv6 address can be divided into three parts:
[IPv6 network prefix] + [subnet segments] + [host address]
- The subnet segments can be split into multiple 4-bit blocks called nibbles
- Each subnet segments represent different geographical or organizational parts of the network. They are represented by 1 or more nibbles.

Example of a prefix:

2001:db8:1234:0000::/64

Section	Description
The yellow highlighted characters	Prefix (48 bits)
The green highlighted characters (zeros)	Place holder for the subnet segments (16 bits)
The red highlighted characters	Subnet mask

The 16 bits that make up the subnet segments can be more granular.

Example: 0011 1111 0000 1101

Segment	Binary	Hexadecimal
Site	0011	0x3
Subsite	1111	0xf
Subnet	0000 1101	0x0d

The resulting network portion of the address is:

2001:db8:1234:3f0d::/64

By changing the mask, the subnet segment could be increased.

2001:db8:1234:0000 0000::/48

2001:db8:1234:0000 0000 0000::/32

This makes more options available for the configuration of the subnet segments. Below is an example of a very basic template:

Edit IPv6 Address Template

Name:

IPv6 Address Prefix:

Subnet Segments ?

Segment Name	Bits	Exclusive	Defined Values
country	4	Disable	
state	4	Disable	
city	4	Disable	
site	4	Disable	
lan	4	Disable	
vlan	4	Disable	

OK Cancel

Using that template, you can see how the GUI could be used to quickly create address objects.

New Address

Category: Address **IPv6 Address** Multicast Address Proxy Address

Name:

Color: Change

Type: IPv6 Template

IPv6 Address Template: test

Subnet Prefix:

Segment Name: country (4 bits): Any Specify

Segment Name: state (4 bits): Any Specify

Segment Name: city (4 bits): Any Specify

Segment Name: site (4 bits): Any Specify

Segment Name: lan (4 bits): Any Specify

Segment Name: vlan (4 bits): Any Specify

Host Type: Any Specify

Host:

Show in Address List: ☒

Comments: 0/255

Tags: Add Tag Category

OK Cancel



You can use the template to enter the subnet prefix alone. You don't have to use the segment portion of the template.

Multicast addresses

Multicast addressing defines a specific range of address values set aside for them. Therefore all IPv4 multicast addresses should be between 224.0.0.0 and 239.255.255.255.

More information on the concepts behind Multicast addressing can be found in the Multicast Forwarding section.

Multicast IP range

This type of address will allow multicast broadcasts to a specified range of addresses.

Creating a multicast IP range address

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**.
 - If you use the down arrow next to **Create New**, select **Address**.
3. Choose the **Category, Multicast Address**
4. Input a **Name** for the address object.
5. Select the **Type, Multicast IP Range** from the drop-down menu.
6. Enter the value for the **Multicast IP Range**
7. Select the **Interface** from the drop-down menu.
8. Enable the **Show in Address List** function
9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example: Multicast IP range address

The company has a large high tech campus that has monitors in many of its meeting rooms. It is common practice for company wide notifications of importance to be done in a streaming video format with the CEO of the company addressing everyone at once.

The video is High Definition quality so takes up a lot of bandwidth. To minimize the impact on the network the network administrators have set things up to allow the use of multicasting to the monitors for these notifications. Now it has to be set up on the FortiGate firewall to allow the traffic.

- The range being used for the multicast is 239.5.5.10 to 239.5.5.200
 - The interface on this FortiGate firewall will be on port 9
1. Go to **Policy & Objects > Objects > Addresses** and select **Create New > Address**.
 2. Fill out the fields with the following information

Category	Multicast Address
Name	Meeting_Room_Displays
Type	Multicast IP Range
Multicast IP Range	239.5.5.10-239.5.5.200
Interface	port9

Show in Address List	<enable>
Comments	<Input into this field is optional>

3. Select **OK**.

4. Enter the following CLI command:

```
config firewall multicast-address
edit "meeting_room_display"
set type multicastrange
set associated-interface "port9"
set start-ip 239.5.5.10
set end-ip 239.5.5.200
set visibility enable
next
end
```

To verify that the address range was added correctly:

1. Go to **Policy & Objects > Objects > Addresses**. Check that the addresses have been added to the address list and that they are correct.
2. Enter the following CLI command:

```
config firewall multicast-address
edit <the name of the address that you wish to verify>
Show full-configuration
```

Broadcast subnet

This type of address will allow multicast broadcast to every node on a subnet.

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**. A drop down menu is displayed. Select **Address**.
3. In the **Category** field, chose **Multicast Address**.
4. Input a **Name** for the address object.
5. In the **Type** field, select **Broadcast Subnet** from the drop down menu.
6. In the **Broadcast Subnet** field enter the address and subnet mask according to the format x.x.x.x/x.x.x.x or the short hand format of x.x.x.x/x. (Remember, it needs to be within the appropriate IP range 224.0.0.0 to 239.255.255.255)
7. In the **Interface** field, leave as the default **any** or select a specific interface from the drop down menu.
8. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
9. Input any additional information in the **Comments** field.
10. Press **OK**.

Example

Field	Value
Category	Broadcast Subnet

Field	Value
Name	Corpnet-B
Type	Broadcast Subnet
Broadcast Subnet	224.5.5.0/24
Interface	any
Show in Address List	[on]
Comments	Corporate Network devices - Broadcast Group B

Multicast IP addresses

Multicast uses the Class D address space. The 224.0.0.0 to 239.255.255.255 IP address range is reserved for multicast groups. The multicast address range applies to multicast groups, not to the originators of multicast packets. The following table lists the reserved multicast address ranges and describes what they are reserved for:

Reserved Multicast address ranges

Reserved Address Range	Use	Notes
224.0.0.0 to 224.0.0.255	Used for network protocols on local networks. For more information, see RFC 1700.	In this range, packets are not forwarded by the router but remain on the local network. They have a Time to Live (TTL) of 1. These addresses are used for communicating routing information.
224.0.1.0 to 238.255.255.255	Global addresses used for multicasting data between organizations and across the Internet. For more information, see RFC 1700.	Some of these addresses are reserved, for example, 224.0.1.1 is used for Network Time Protocol (NTP).
239.0.0.0 to 239.255.255.255	Limited scope addresses used for local groups and organizations. For more information, see RFC 2365.	Routers are configured with filters to prevent multicasts to these addresses from leaving the local system.

Creating multicast security policies requires multicast firewall addresses. You can add multicast firewall addresses by going to **Firewall Objects > Address > Addresses** and selecting **Create New > Multicast Address**. The factory default configuration includes multicast addresses for Bonjour (224.0.0.251-224.0.0.251), EIGRP (224.0.0.10-224.0.0.100), OSPF (224.0.0.5-224.0.0.60), all_hosts (224.0.0.1-224.0.0.1), and all_routers (224.0.0.2-224.0.0.2).

Proxy addresses

This category of address is different from the other addresses in that it is not designed to be used in the normal firewall policy configuration. It is intended to be used only with explicit web proxies.

In some respects they can be like a FQDN addresses in that they refer to an alpha-numeric string that is assigned to an IP address, but then goes an additional level of granularity by using additional information and criteria to further specify locations or types of traffic within the website itself. In depth information on Explicit Proxy Addressing can be found in [WAN Optimization](#), but it is worth laying out the steps of how to create an address object for this category.

Creating an proxy address

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**. A drop down menu is displayed. Select **Address**.
3. In the **Category** field, chose **Proxy Address**.
4. Input a **Name** for the address object.
5. For the **Type** field, select one of the options from the drop down menu.

Within the Explicit Proxy Address category there are 8 types of addresses. Each of these types will have associated field(s) that also need to have values entered to make the object specific to it's address.

Type = URL Pattern

- In the **Host** field, choose from drop down menu
- In the **URL Path Regex** field, enter the appropriate string

Host Regex Match

- In the **Host Regex Pattern** field, enter the appropriate string

URL Category

- In the **Host** field, choose from drop down menu
- In the **URL Category** field, choose from drop down menu

HTTP Method

- In the **Host** field, choose from drop down menu
- In the **Request Method** field, choose from drop down menu

The options are:

- CONNECT
- DELETE
- GET
- HEAD
- OPTIONS
- POST
- PUT
- TRACE

User Agent

- In the **Host** field, choose from drop down menu
- In the **User Agent** field, choose from drop down menu

The options are:

- Apple Safari
- Google Chrome
- Microsoft Internet Explorer or Spartan
- Mozilla Firefox
- Other browsers

HTTP Header

- In the **Host** field, choose from drop down menu
- In the **Header Name** field, enter the appropriate string value
- In the **Header Regex** field, enter the appropriate string value

Advanced (Source)

- In the **Host** field, choose from drop down menu
- In the **Request Method** field, choose from drop down menu (see **HTTP Method** type for option list)
- In the **User Agent** field, choose from drop down menu (see **User Agent** type for option list)
- In the **Header Group** table, create, edit or delete **Header Name** strings and associated **Header Regex** strings

Advance (Destination)

- In the **Host** field, choose from drop down menu
- In the **Host Regex Pattern** field, enter the appropriate string
- In the **URL Category** field, choose from drop down menu

6. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled the address will appear in drop down menus where it is an option.
7. Input any additional information in the **Comments** field.
8. Press **OK**.

Proxy address groups

To create a Proxy address group:

1. Go to **Policy & Objects > Addresses**.
2. Click on **+ Create New** to get the drop down menu. Select **Address Group**.
3. In the **Category** field, choose **Proxy Group**.
4. Fill in a descriptive name in the **Group Name** field.
5. If you wish, use the **Change** link to change the **Color** of icons in the GUI. There are 32 color options.
6. In the Type field, select whether the group will be a **Source Group** (composed of source addresses) or a **Destination Group** (composed of destination addresses).
7. Select anywhere in the **Members** field to bring forth the pane of potential members for selection to the group.
8. Select the desired on/off toggle setting for **Show in Address List**. If the setting is enabled, the address will appear in drop down menus where it is an option.
9. Input any additional information in the **Comments** field.
10. Click on **OK**.

New Address Group

Category

IPv4 Group

IPv6 Group

Proxy Group

Group Name

Color

[Change]

Type

Source Group

Destination Group

Members

+

Show in Address List

☒

Comments

Write a comment...

0/255

OK

Cancel

Internet services

In FortiOS 5.4, support was added for Internet Service objects which could be used with **FortiView**, **Logging**, **Routing** and **WAN Load Balancing**. Now they can be added to firewall policies as well.



There is an either or relationship between Internet Service objects and destination address and service combinations in firewall policies. This means that a destination address and service can be specified in the policy OR an Internet service, not both.

CLI

The related CLI options/syntax are:

```
config firewall policy
edit 1
set internet-service 1 5 10
set internet-service-custom test
set internet-service-negate [enable|disable]
end
```

GUI

In the policy listing page you will notice that if an Internet Service object is used, it will be found in both the **Destination** and **Service** column.

In the policy editing page the **Destination Address**, now **Destination** field now has two types, **Address** and **Internet Service**.

New Policy		Select Entries
Name	Citrix access	Address Internet Service
Incoming Interface	port1	Search
Outgoing Interface	port2	<ul style="list-style-type: none"> Citrix-FTP(S) Citrix-IMAP(S) Citrix-NetBIOS.Name.Service Citrix-NetBIOS.Session.Service Citrix-SMTP(S) Citrix-SSH Citrix-Web CNN-FTP(S) CNN-SMTP(S) Dropbox-DNS Dropbox-NetBIOS.Name.Service
Source	all	
Destination	<ul style="list-style-type: none"> Citrix-DNS Citrix-FTP(S) Citrix-IMAP(S) Citrix-NetBIOS.Name.Service Citrix-NetBIOS.Session.Service Citrix-SMTP(S) Citrix-SSH Citrix-Web CNN-FTP(S) 	

Address groups

Address groups are designed for ease of use in the administration of the device. If you have a number of addresses or address ranges that will commonly be treated the same or require the same security policies, you can put them into address groups, rather than entering multiple individual addresses in each policy refers to them.

The use of groups is not required. If you have a number of different addresses you could add them individually to a policy and the FortiGate firewall will process them just as quickly and efficiently as if they were in a group, but the chances are that if you have used a group once you could need to use it again and depending on the number of addresses involved entering them individually for each policy can become tedious and the likelihood of an address being missed becomes greater. If you have a number of policies using that combination of addresses it is much easier to add or subtract addresses from the group than to try and remember all of the firewall policies that combination of addresses was used in. With the group, you only have to make the one edit and it is used by any firewall policy using that address group.

Because security policies require addresses with homogenous network interfaces, address groups should contain only addresses bound to the same network interface, or to Any.

For example, if address 1.1.1.1 is associated with port1, and address 2.2.2.2 is associated with port2, they cannot be in the same group. However, if 1.1.1.1 and 2.2.2.2 are configured with an interface of Any, they can be grouped, even if the addresses involve different networks.

There are 3 Categories of Address groups to choose from:

- IPv4 Group
- IPv6 Group
- Proxy Group

You cannot mix different categories of addresses within a group, so whether or not it makes sense from an administrative purpose to group certain addresses together, if some are IPv4 and some are IPv6, it cannot be done.

Creating an address group

1. Go to **Policy & Objects > Addresses**.
2. Select the down arrow next to **Create New**, select **Address Group**.
3. Choose the **Category**, that is applicable to the proposed selection of addresses.
4. Input a **Group Name** for the address object.

Depending on which **Category** has been chosen the configurations will differ slightly

IPv4 group

1. Select the "+" in the **Members** field. You can select members of the group from the window that slides out from the left of the screen. It is possible to select more than 1 entry. Select the "X" icon in the field to remove an entry.
2. Select the desired on/off toggle setting for **Show in Address List**.
3. Select the desired on/off toggle setting for **Static Route Configuration**.

IPv6 group

1. Select the "+" in the **Members** field. You can select members of the group from the window that slides out from the left of the screen. It is possible to select more than 1 entry. Select the "X" icon in the field to remove an entry.
2. Select the desired on/off toggle setting for **Show in Address List**.

Proxy group

1. Select which Type, either **Source Group** or **Destination Group**.
2. Select the "+" in the **Members** field. You can select members of the group from the window that slides out from the left of the screen. It is possible to select more than 1 entry. Select the "X" icon in the field to remove an entry.
3. Select the desired on/off toggle setting for **Show in Address List**.

Irrespective of the Category the groups all have the same final configuration options:

1. Input any additional information in the **Comments** field.
2. Press **OK**.

UUID support

Syntax:

```
config firewall {address|addres6|addgrp|addgrp6}
edit 1
    set uuid <example uuid: 8289ef80-f879-51e2-20dd-fa62c5c51f44>
next
end
```

Virtual IPs

The mapping of a specific IP address to another specific IP address is usually referred to as Destination NAT. When the Central NAT Table is not being used, FortiOS calls this a Virtual IP Address, sometimes referred to as a VIP. FortiOS uses a DNAT or Virtual IP address to map an External IP address to an IP address. This address does not have to be an individual host, it can also be an address range. This mapping can include all TCP/UDP

ports or if Port Forwarding is enabled it will only refer to the specific ports configured. Because, the Central NAT table is disabled by default the term Virtual IP address or VIP will be used predominantly.

Virtual IP addresses are typically used to NAT external or Public IP addresses to internal or Private IP addresses. Using a Virtual IP address between 2 internal Interfaces made up of Private IP addresses is possible but there is rarely a reason to do so as the 2 networks can just use the IP addresses of the networks without the need for any address translation. Using a Virtual IP address for traffic going from the inside to the Internet is even less likely to be a requirement, but it is supported.

Something that needs to be considered when there are multiple Public IP addresses on the external interface(s) is that when a Virtual IP address is used without Port Forwarding enabled there is a reciprocal effect as far as traffic flow is concerned. Normally, on a firewall policy where NAT is enabled, for outgoing traffic the internal address is translated to the Public address that is assigned to the FortiGate, but if there is a Virtual IP address with no port forwarding enabled, then the Internal IP address in the Mapped field would be translated to the IP address configured as the External Address in the VIP settings.

Example

- The assigned External address (WAN1) of the FortiGate unit is 172.12.96.3 with a subnet mask of 255.255.255.128
- There is a Virtual IP address set up to map the external address 172.12.96.127 on WAN1 to the internal IP address of 192.168.1.127
- Port Forwarding is not enabled because you want all allowed traffic going to the external IP address to go to this server.

In this case any outbound traffic from 192.168.1.127 will go out on WAN1 with the IP address of 172.12.96.127 as the source IP address.

In terms of actually using the Virtual IP address, they would be using in the security policies in the same places that other addresses would be used, usually as a Destination Address.

UUID support for VIP

UUID is now supported in for virtual IPs and virtual IP groups. This includes virtual IPs for IPv4, IPv6, NAT46, and NAT64. To view the UUID for these objects in a FortiGate unit's logs, log-uuid must be set to extended mode, rather than policy-only (which only shows the policy UUID in a traffic log). UUID can only be configured through the CLI

Syntax

```
config sys global
  set log-uuid {disable | policy-only | extended}
end
```



There is another type of address that the term “virtual IP address” commonly refers to which is used in load balancing and other similar configurations. In those cases, a number of devices share a separately created virtual IP address that can be sent to multiple possible devices. In FortiOS these are referred to as Virtual Servers and are configured in the “Load Balance” section.



If Central-NAT is enabled in the CLI the GUI will be different.

Instead of **VIP Type**, the field label will be **DNAT & VIP Type**

Instead of **IPv4** the option will be **IPv4 DNAT**

There will also be the addition setting of **Source Interface Filter**.

Commands to set `central-nat`:

```
config system settings
    set central-nat [enable | disable]
end
```

Creating a virtual IP

1. Go to **Policy & Objects > Virtual IPs**.
2. Select **Create New**. A drop down menu is displayed. Select **Virtual IP**.
3. From the **VIP Type** options, choose an applicable type based on the IP addressing involved. Which is chosen will depend on which of the IP version networks is on the external interface of the FortiGate unit and which is on the internal interface.

The available options are:

- **IPv4** - IPv4 on both sides of the FortiGate Unit.
 - **IPv6** - IPv6 on both sides of the FortiGate Unit.
 - **NAT46** - Going from an IPv4 Network to an IPv6 Network.
 - **NAT64** - Going from an IPv6 Network to an IPv4 Network.
4. In the **Name** field, input a unique identifier for the Virtual IP.
 5. Input any additional information in the **Comments** field.
 6. The **Color** of the icons that represent the object in the GUI can be changed by clicking on the **[Change]** link and choosing from the 32 colors.

Because the configuration differs slightly for each type the next steps will be under a separate heading based on the type of the VIP

Configuring a VIP for IPv4

In the **Network** section:

7. If an IPv4 type of Virtual IP, select the **Interface** setting.
Using the drop down menu for the Interface Field, choose the incoming interface for the traffic.
The IPv4 VIP Type is the only one that uses this field. This is a legacy function from previous versions so that they can be upgraded without complicated reconfiguration. The External IP address, which is a required field, tells the unit which interface to use so it is perfectly acceptable to choose **"any"** as the interface. In some configurations, if the Interface field is not set to **"any"** the Virtual IP object will not one of the displayed options when choosing a destination address.
8. Configure the **External IP Address/Range**.

There are two fields. If there is a single IP address, use that address in both fields. This will be the address on the outside of the network that is usually the public address of the server. The format of the address will depend on the **VIP Type** option that was selected.

9. Configure the **Mapped IP Address/Range**. This will be the address that the traffic is being directed to. There are two fields. If there is a single IP address, use that address in both fields. The format of the address will depend on the **VIP Type** option that was selected.

In the **Optional Filters**

10. Disable/Enable the **Optional Filters**.
If only specific IP addresses and/or services are allowed to be the source for traffic using the VIP, enable the **Optional Filters**.
11. To specify an allowed address enter the value in the field labeled **Source Address**. The value can be formatted in three different ways.
 - **Source IP** - Use the standard format for a single IP address
 - **Range** - Enter the first and last members of the range
 - **Subnet** - Enter the IP address of the broadcast address for the subnet.
To add additional addresses, click on the "+" below the last field with an address. To subtract an address, click on the "X" next to the field you wish to delete.
12. To specify an allowed Service, toggle the **Services** option to enabled. Set the **Services** parameter by selecting the field with the "+" in the field. This will slide a window out from the right. Single or multiple options can be selected by highlighting the services wanted, unless the **ALL** option is chosen, in which case it will be the only option. For more information on services, check the Firewall Objects section called [Services and TCP ports](#).
13. Disable/Enable **Port Forwarding**. If only the traffic for a specific port or port range is being forwarded, enable this setting.
14. Select the **Protocol** from
 - **TCP**
 - **UDP**
 - **SCTP**
 - **ICMP**
15. Configure the **External Service Port**. This is the port(s) on the external interface of the FortiGate (the destination port in the header of the packets). The first field is for the first port in the range the second is for the last port in the range. As you enter a value in the first field, the second field will auto populate with the same number, working on the premise that a single port is common. Just edit the second field to extend the range.
16. Configure the setting **Map to Port**. This will be the listening port on the device located on the internal side of the network. It does not have to be the same as the **External Service Port**. The first field is for the first port in the range the second is for the last port in the range. As you enter a value in the first field, the second field will auto populate with the same number, working on the premise that a single port is common. Just edit the second field to extend the range.
17. Press **OK**.

Example

This example is for a VIP that is being used to direct traffic from the external IP address to a web server on the internal network. The web server is for company use only. The company's public facing web server already used

port 80 and there is only one IP external IP address so the traffic for this server is being listened for on port 8080 of the external interface and being sent to port 80 on the internal host.

Field	Value
VIP Type	IPv4
Name	Internal_Webserver
Comments	Web server with Collaboration tools for Corporate employees
Interface	Any
External IP Address/Range	172.13.100.27 <this would normally be a public IP address>
Mapped IP Address/Range	192.168.34.150
Optional Filters	enabled
Source Address Filter	<list of IP addresses of remote users>
Services	enabled with HTTP in the list
Port Forwarding	enabled
Map to Port	80 - 80

Configuring a VIP for IPv6

In the **Network** section:

7. Configure the **External IP Address/Range**.

There are two fields. If there is a single IP address, use that address in both fields. This will be the address on the outside of the network that is usually the public address of the server. Enter the address in the standard IPv6 format.

8. Configure the **Mapped IP Address/Range**. This will be the address that the traffic is being directed to.

There are two fields. If there is a single IP address, use that address in both fields. Enter the address in the standard IPv6 format.

In the **Optional Filters**

9. Disable/Enable the **Optional Filters**.

If only specific IP addresses and/or services are allowed to be the source for traffic using the VIP, enable the **Optional Filters**.

10. To specify an allowed address enter the value in the field labeled **Source Address**. The value can be formatted in three different ways.

- **Source IP** - Use the standard format for a single IP address
 - **Range** - Enter the first and last members of the range
 - **Subnet** - Enter the IP address of the broadcast address for the subnet.
To add additional addresses, click on the "+" below the last field with an address. To subtract an address, click on the "X" next to the field you wish to delete.
12. Disable/Enable **Port Forwarding**. If only the traffic for a specific port or port range is being forwarded, enable this setting.
 13. Select the **Protocol** from
 - **TCP**
 - **UDP**
 - **SCTP**
 14. Configure the **External Service Port**. This is the port(s) on the external interface of the FortiGate (the destination port in the header of the packets). The first field is for the first port in the range the second is for the last port in the range. As you enter a value in the first field, the second field will auto populate with the same number, working on the premise that a single port is common. Just edit the second field to extend the range.
 15. Configure the setting **Map to Port**. This will be the listening port on the device located on the internal side of the network. It does not have to be the same as the **External Service Port**. The first field is for the first port in the range the second is for the last port in the range. As you enter a value in the first field, the second field will auto populate with the same number, working on the premise that a single port is common. Just edit the second field to extend the range.
 16. Press **OK**.

Configuring a VIP for NAT46

In the **Network** section:

7. Configure the **External IP Address/Range**.
There are two fields. If there is a single IP address, use that address in both fields. This will be the address on the outside of the network that is usually the public address of the server. Enter the address in the standard IPv4 format.
8. Configure the **Mapped IP Address/Range**. This will be the address that the traffic is being directed to.
There are two fields. If there is a single IP address, use that address in both fields. Enter the address in the standard IPv6 format.

In the **Optional Filters**

9. Disable/Enable the **Optional Filters**.
If only specific IP addresses and/or services are allowed to be the source for traffic using the VIP, enable the **Optional Filters**.
10. To specify an allowed address enter the value in the field labeled **Source Address**. The value can be formatted in three different ways.
 - **Source IP** - Use the standard format for a single IP address
 - **Range** - Enter the first and last members of the range
 - **Subnet** - Enter the IP address of the broadcast address for the subnet.

To add additional addresses, click on the "+" below the last field with an address. To subtract an address, click on the "X" next to the field you wish to delete.

12. Disable/Enable **Port Forwarding**. If only the traffic for a specific port or port range is being forwarded, enable this setting.
13. Select the **Protocol** from
 - **TCP**
 - **UDP**
14. Configure the **External Service Port**. This is the port(s) on the external interface of the FortiGate (the destination port in the header of the packets). The first field is for the first port in the range the second is for the last port in the range. As you enter a value in the first field, the second field will auto populate with the same number, working on the premise that a single port is common. Just edit the second field to extend the range.
15. Configure the setting **Map to Port**. This will be the listening port on the device located on the internal side of the network. It does not have to be the same as the **External Service Port**. The first field is for the first port in the range the second is for the last port in the range. As you enter a value in the first field, the second field will auto populate with the same number, working on the premise that a single port is common. Just edit the second field to extend the range.
16. Press **OK**.



In order for VIP46 to work correctly, NAT64 must be enabled. To enable NAT64, set the following configuration:

```
conf sys nat64
  set status enable
end
```

Configuring a VIP for NAT64

In the **Network** section:

7. Configure the **External IP Address/Range**.
There are two fields. If there is a single IP address, use that address in both fields. This will be the address on the outside of the network that is usually the public address of the server. Enter the address in the standard IPv6 format.
8. Configure the **Mapped IP Address/Range**. This will be the address that the traffic is being directed to.
There are two fields. If there is a single IP address, use that address in both fields. Enter the address in the standard IPv4 format.

In the **Optional Filters**

9. Disable/Enable the **Optional Filters**.
If only specific IP addresses and/or services are allowed to be the source for traffic using the VIP, enable the **Optional Filters**.
10. To specify an allowed address enter the value in the field labeled **Source Address**. The value can be formatted in three different ways.
 - **Source IP** - Use the standard format for a single IP address
 - **Range** - Enter the first and last members of the range
 - **Subnet** - Enter the IP address of the broadcast address for the subnet.

To add additional addresses, click on the "+" below the last field with an address. To subtract an address, click on the "X" next to the field you wish to delete.

12. Disable/Enable **Port Forwarding**. If only the traffic for a specific port or port range is being forwarded, enable this setting.
13. Select the **Protocol** from
 - **TCP**
 - **UDP**
14. Configure the **External Service Port**. This is the port(s) on the external interface of the FortiGate (the destination port in the header of the packets). The first field is for the first port in the range the second is for the last port in the range. As you enter a value in the first field, the second field will auto populate with the same number, working on the premise that a single port is common. Just edit the second field to extend the range.
15. Configure the setting **Map to Port**. This will be the listening port on the device located on the internal side of the network. It does not have to be the same as the **External Service Port**. The first field is for the first port in the range the second is for the last port in the range. As you enter a value in the first field, the second field will auto populate with the same number, working on the premise that a single port is common. Just edit the second field to extend the range.
16. Press **OK**.

FQDN in VIPs

Instead of mapping to an IP address a VIP can use a FQDN(Fully Qualified Domain Name). This has to be configured in the CLI and the FQDN must be an address object that is already configured in the address listing.

The syntax for using a FQDN is:

```
config firewall vip
  edit <VIP id>
    set type fqdn
    set mapped-addr <FQDN address object>
  end
```

Dynamic VIP according to DNS translation

When a dynamic virtual IP is used in a policy, the dynamic DNS translation table is installed along with the dynamic NAT translation table into the kernel. All matched DNS responses will be translated and recorded regardless if they hit the policy. When a client request hits the policy, dynamic NAT translation will occur if it matches a record, otherwise the traffic will be blocked.

Syntax

```
config firewall vip
  edit "1"
    set type dns-translation
    set extip 192.168.0.1-192.168.0.100
    set extintf "dmz"
    set dns-mapping-ttl 604800
    set mappedip "3.3.3.0/24" "4.0.0.0/24"
  end
end
```

Virtual IP groups

Just like other address, Virtual IP addresses can be organized into groups for ease of administration. If you have multiple virtual IPs that are likely to be associated to common firewall policies rather than add them individually to each of the policies you can add the instead. That way, if the members of the group change then any changes made to the group will propagate to all of the policies using that group.

When using a Virtual IP address group the firewall policy will take into account all of the configured parameters of the Virtual IPs: IP addresses, Ports and port types.

Creating a virtual IP group

1. Go to **Policy & Objects > Virtual IPs**.
2. Select **Create New**. A drop down menu is displayed. Select **Virtual IP Group**.
3. Select the **Type** for VIP group you wish to create.

The options available are:

- **IPv4** - IPv4 on both sides of the FortiGate Unit.
- **IPv6** - IPv6 on both sides of the FortiGate Unit.
- **NAT46** - Going from an IPv4 Network to an IPv6 Network.
- **NAT64** - Going from an IPv6 Network to an IPv4 Network.

Which is chosen will depend on which of the IP version networks is on the external interface of the FortiGate unit and which is on the internal interface. The options will be:

4. Enter a unique identifier for the group in the **Name** field.
5. Enter any additional information in the **Comments** field.
6. If you wish, use the **Change** link to change the **Color** of icons in the GUI. There are 32 color options.
7. If the **Type** is **IPv4**, the **Interface** field will be available. Use the drop-down menu to select the interface if all of the VIPs are on the same interface. If any of the VIPs are on different interfaces or if any of them are associated with the "any" option, choose the any option for the group.
8. Select anywhere in the **Members** field to bring forth the pane of potential members for selection to the group.
9. Press **OK**.

Configuring IP pools

An IP pool is essentially one in which the IP address that is assigned to the sending computer is not known until the session is created, therefore at the very least it will have to be a pool of at least 2 potential addresses. A quick example would be an IP pool for users of a VPN. IP pools are based upon the version of IP determined by the interface that they are associated with so as expected there are two types of IP pools that can be configured:

- ["Creating a IPv4 pool" on page 736](#)
- ["Creating a IPv6 pool" on page 741](#)

Because of the differences in the configuration for the two types of pools, instructions for configuring them will be done separately.

Creating a IPv4 pool

1. Go to **Policy & Objects > IP Pools**.
2. Select **Create New**.
3. In the **IP Pool Type** field choose **IPv4 Pool**

4. Enter a name in the **Name** field for the new service
5. Include any description you would like in the **Comments** field
6. In the **Type** field choose between:
 - **Overload**
 - **One-to-One**
 - **Fixed Port Range**
 - **Port Block Allocation**

At this point the configurations can start to differ based on the type of type of pool.

For more information on the different types of IP pools, check [IP Pools](#) in the Concepts section.

Overload

7. For the **External IP Range** fields, enter the lowest and highest addresses in the range. If you only want a single address used, enter the same address in both fields.
8. Enable the **ARP Reply** field by making sure there is a check in the box
9. Select **OK**

Overload example for GUI

In this example, the Sales team needs to connect to an Application Service Provider that does the accounting for the company. As a security measure, the ASP only accepts traffic from a white list of IP addresses. There is 1 public IP address of the company on that list. The Sales team consists of 40 people, so they need to share. The external interface is wan1.

Field	Value
IP Pool Type	IPv4 Pool
Name	Sales_Team
Comments	For the Sales team to use to connect to the Accounting ASP
Type	Overload (This is the default)
External IP Range	10.23.56.20 - 10.23.56.20
ARP Reply	enabled

Overload example for CLI

```
config firewall ippool
  edit Sales_Team
    set comments "For the Sales team to use to connect to the Accounting ASP"
    set type overload
    set startip 10.23.56.20
    set endip 10.23.56.20
    set arp-reply enable
    set arp-intf wan1
  end
```


One-to-one

7. For the **External IP Range** fields, enter the lowest and highest addresses in the range. If you only want a single address used, enter the same address in both fields.
8. Enable the **ARP Reply** field by making sure there is a check in the box.
9. Select **OK**

One-to-one example for GUI

In this example, the external IP address of the mail server is part of a range assigned to the company but not the one that is assigned to the Internet facing interface. A VIP has been set up but in order to properly resolve Reverse DNS lookups the mail server always has to use a specific IP address. The external interface is wan1.

Field	Value
IP Pool Type	IPv4 Pool
Name	Mail-Server
Comments	So the correct IP address is resolved on Reverse DNS look ups of the mail server.
Type	One-to-one
External IP Range	10.23.56.21 - 10.23.56.21
ARP Reply	enabled

One-to-one example for CLI

```
config firewall ippool
  edit Mail-Server
    set comments "So the the correct IP address is resolved on reverse DNS look ups of
      the mail server."
    set type one-to-one
    set startip 10.23.56.21
    set endip 10.23.56.21
    set arp-reply enable
    set arp-intf wan1
  end
```

Fixed port range

7. For the **External IP Range** fields, enter the lowest and highest addresses in the range. If you only want a single address used, enter the same address in both fields.
8. Fort the **Internal IP Range** fields, enter the lowest and highest addresses in the range.
9. Enable the **ARP Reply** field by making sure there is a check in the box
10. Select **OK**

Fixed port range example for GUI

In this example, the company has a range of 10 IP address that they want to be used by employees on a specific subnet for NATing. The external interface is wan1.

Field	Value
IP Pool Type	IPv4 Pool
Name	IPPool-3
Comments	IP range to be used by outgoing traffic
Type	Fixed Port Range
External IP Range	10.23.56.22 - 10.23.56.31
Internal IP Range	192.168.23.1 - 192.168.23.254
ARP Reply	enabled

Fixed port range example for CLI

```
config firewall ippool
edit IPPool-3
    set comments "So the the correct IP address is resolved on reverse DNS look ups of
        the mail server."
    set type fixed-port-range
    set startip 10.23.56.22
    set endip 10.23.56.31
    set source-startip 192.168.23.1
    set source-endip 192.168.23.254
    set arp-reply enable
    set arp-intf wan1
end
```

Port block allocation

- For the **External IP Range** fields, enter the lowest and highest addresses in the range. If you only want a single address used, enter the same address in both fields.
- In the **Block Size** field, either type in the value or use the up or down arrows to set the value of the block size.
- In the **Blocks Per User** field, either type in the value or use the up or down arrows to set the value for the number of blocks per user.
- Enable the **ARP Reply** field by making sure there is a check in the box
- Select **OK**

Port block allocation timeout

The port block allocation timeout value is configurable. The setting is found in the CLI.

The option `pba-timeout` has been added to the firewall ip pool configuration. The availability of this option is dependent on the type option being set to port-block-allocation. The timeout value is measured in seconds and is an integer between 3 and 300, with the default being 30.

Syntax:

```
config firewall ippool
  edit <name of PBA pool>
    set type port-block-allocation
    set pba-timeout <integer>
  end
```

Port block allocation example for GUI

In this example, an small ISP is setting up NATing for its clients, but to be fair it is putting some restrictions on the number of connections each client can have so that no one hogs all of the possible ports and addresses. The external interface is port12.

Field	Value
IP Pool Type	IPv4 Pool
Name	Client-IPPool
Comments	IP Pool for clients to access the Internet
Type	Port Block Allocation
External IP Range	10.23.75.5 - 10.23.75.200
Block Size	64
Blocks Per User	8
ARP Reply	enabled

Port block allocation example for CLI

```
config firewall ippool
  edit Client-IPPool
    set comments "IP Pool for clients to access the Internet"
    set type port-block-allocation
    set startip 10.23.75.5
    set endip 10.23.75.200
    set block-size 64
    set num-blocks-per-user 8
    set permit-any-host disable
    set arp-intf wan1
    set arp-reply enable
    set arp-intf port12
  end
```

Creating a IPv6 pool

1. Go to **Policy & Objects > IP Pools**.
2. Select **Create New**.
3. In the IP Pool Type field choose **IPv6 Pool**
4. Enter a name in the **Name** field for the new service
5. Include any description you would like in the **Comments** field
6. For the **External IP Range** fields, enter the lowest and highest addresses in the range.

IPv6 example for GUI

In this example, there is a similar situation to the One-to-one example earlier. There is a mail server that needs to be resolved to a specific IP address in Reverse DNS look-ups. The difference in this case is the company is an early adopter of IPv6 connectivity to the Internet.

Field	Value
IP Pool Type	IPv6 Pool
Name	Mail-svr-ipv6
Comments	Registered IPv6 address for mail server
External IP Range	fd2f:50ec:cdea:0663::1025 - fd2f:50ec:cdea:0663::1025

Port block allocation example for CLI

```
config firewall ippool6
edit Mail-svr-ipv6
set comments "Registered IPv6 address for mail server"
set startip fd2f:50ec:cdea:663::102
set endip fd2f:50ec:cdea:663::1025
end
```

Creating NAT46 IP pool and multiple (secondary) NAT64 prefixes

Policies that translate between IPv4 and IPv6 can use IPv4 address pools or IPv6 prefixes to be used in the policies, giving more options to the configuration of addresses.

NAT46

For using the ippool in NAT46 policies, first enable the use of ippools and then set the names of the ippool(s).

```
config firewall policy46
edit 1
set uuid e9c6ca3e-72ea-51e7-554a-1185693d03eb
set srcintf "wan1"
set dstintf "internal7"
set srcaddr "external-net4"
set dstaddr "internal-vip46"
set action accept
set schedule "always"
set service "ALL"
set ippool enable
```

```
set poolname "intit-pool6"  
end
```

NAT64

In order to use these options in the NAT64 firewall policies the new settings `secondary-prefix status` and `secondary-prefix options` have to be configured as in the example below.

```
config system nat64  
set nat64-prefix 2001::/96  
set secondary-prefix enable  
config secondary-prefix  
edit 1  
set nat64-prefix 2002::/94  
next  
edit 2  
set nat64-prefix 2003::/95  
end  
end
```



The primary prefix must have a length of 96, but the secondary prefixes can be different lengths

Services

While there are a number of services already configured within FortiOS, the firmware allows for administrators to configure their own. The reasons for doing this usually fall into one or more of the following categories:

- The service is not common enough to have a standard configuration
- The service is not established enough to have a standard configuration
- The service has a standard port number but there is a reason to use a different one:
 - Port is already in use by another service
 - For security reasons, want to avoid standard port

When looking at the list of preconfigured services it may seem like there are a lot, but keep in mind that the theoretical limit for port numbers is 65,535. This gives a fairly good sized range when you are choosing what port to assign a service but there are a few points to keep in mind.

- Most of the well known ports are in the range 0 - 1023
- Most ports assigned by the Internet Corporation for Assigned Names and Numbers (ICANN) will be in the 1024 - 49151 range
- Port numbers between 49,152 and 65,535 are often used for dynamic, private or ephemeral ports.

There are 3 Service objects that can be added and configured:

- Categories
- Services
- Service Groups

Categories

In order to make sorting through the services easier, there is a field to categorize the services. Because selecting a category is part of the process for creating a new service, the configuration of categories will be explained first.

The services can be sorted into the following groups:

- General
- Web Access
- File Access
- Email
- Network Services
- Authentication
- Remote Access
- Tunneling
- VoIP, Messaging and Other Applications
- Web Proxy
- Uncategorized

The categories are for organization purposes so there is not many settings when creating a new one.

Creating a new service category

1. Go to **Policy & Objects > Services**.
2. Select **Create New**. A drop down menu is displayed. Select **Category**
3. Input a **Name** for the category.
4. Input any additional information in the **Comments** field.
5. Press **OK**.

Example

You plan on adding a number of devices such as web cameras that will allow the monitoring of the physical security of your datacenter. A number of non-standard services will have to be created and you would like to keep them grouped together under the heading of "Surveillance"

Example of a new category in the GUI

1. Go to **Policy & Objects > Objects > Services** and select **Create New > Category**.
2. Fill out the fields with the following information

Field	Value
Name	Surveillance
Comments	For DataCenter Surveillance Devices

3. Select **OK**.

Example of a New Category in the CLI

Enter the following CLI command:

```
config firewall service category
edit Surveillance
set comment "For DataCenter Surveillance Devices"
end
```

To verify that the category was added correctly:

1. Go to **Policy & Objects > Objects > Services**. Select the Category Settings icon . A listing of the categories should be displayed.
2. Enter the following CLI command:

```
config firewall service category
show
```

This should bring up all of the categories. Check to see that the new one is displayed.

Configuring a new service

Occasionally, the preconfigured list of services will not contain the needed service. There are a few variations in the creation of a service depending upon the protocol type, but the first steps in the creation of the service are common to all the variations.

To create a new service:

1. Go to **Policy & Objects > Services**.
2. Select **Create New**. A drop down menu is displayed. Select **Service**
3. Enter a name in the **Name** field for the new service
4. Include any description you would like in the **Comments** field
5. In the **Service Type** field choose between **Firewall** and **Explicit Proxy**.
6. Enable the toggle in the **Show in Service List**. If you can't see the service when you need to select it, it serves very little purpose.
7. For the **Category** field, choose the appropriate category from the **Category** drop down menu. If none is chosen, the **Uncategorized** option will be chosen by default.

Protocol options

This is the section where the configuration options of the service will differ depending on the type of protocol chosen. (The Step numbers will all continue on from the common step sequence).

The protocol options for **Firewall** service type are:

- **TCP/UDP/SCTP**
- **ICMP**
- **ICMP6**
- **IP**

The protocol options for **Proxy** service type are:

- **ALL**
- **CONNECT**
- **FTP**
- **HTTP**

- **SOCKS-TCP**
- **SOCKS-UDP**

TCP/UDP/SCTP

- For the **Protocol Type** field, choose **TCP/UDP/SCTP** from the drop down menu
- For the **Address** field, choose IP Range or **FQDN** (Fully Qualified Domain Name) if there is to be a specific destination for the service. Depending on which type of address is selected, the field value needs to be filled with a FQDN string or an IP address in one of the 3 standard IPv4 address formats:
 - x.x.x.x - for a specific address
 - x.x.x.x/x - for a subnet
 - x.x.x.x-x.x.x.x - for a range of specific addresses
- Configure the **Destination Port** by:
 - Select from the drop down menu, **TCP**, **UDP** or **SCTP**
 - Enter the low end to the port range in the field indicated by grayed out **Low**.
 - Enter the high end of the port range in the field indicated by grayed out **High**. If there is only a single port in the range **High** can be left empty
 - Multiple ports or port ranges can be added by using the "+" at the beginning of the row
 - Rows can be removed by using the trash can symbol at the end of the row
- If required, you can **Specify Source Ports** for the service by enabling the toggle switch.
 - The **Src Port** will match up with a **Destination Port**
 - **Src Ports** cannot be configured without there being a value for the **Destination Port**
 - The same rules for configuring the **Destination Ports** applies to the **Src Ports**
- Select **OK** to confirm the configuration

Example

Example settings for a TCP protocol service. In this case, it is for an administrative connection to web servers on the DMZ. The protocol used is HTTPS which would normally use port 443, but that is already in use by another service such as Admin access to the firewall or an SSL-VPN connection.

Field	Value
Name	Example.com_WebAdmin
Comments	Admin connection to Example.com Website
Service Type	Firewall
Show in Service List	enabled
Category	Web Access
Protocol Options	
Protocol Type	TCP/UDP/SCTP

Field	Value
IP/FQDN	<left blank>
Destination Port	<ul style="list-style-type: none"> • Protocol: TCP • Low: 4300 • High: <left blank>
Specify Source Ports	<disabled>

Creating a new TCP/UDP/SCTP service in the CLI

The following is the creation of the same service using the command line.

```
config firewall service custom
edit Example.com_WebAdmin
set comment "Admin connection to Example.com Website"
set category Web Access
set protocol TCP/UDP/SCTP
set tcp-portrange 4300
end
end
```

ICMP / ICMP6

- For the **Protocol Type** field, choose **ICMP** or **ICMP6** from the drop down menu
- In the **Type** field enter the appropriate type number based on the information found in ["ICMP Types and Codes" on page 1](#) or in ["ICMPv6 Types and Codes" on page 1](#), depending on whether the **Protocol Type** is **ICMP** or **ICMPv6**
- In the **Code** field enter the appropriate code number for the type, if applicable, based on the information found in ["ICMP Types and Codes" on page 1](#) or in ["ICMPv6 Types and Codes" on page 1](#), depending on whether the **Protocol Type** is **ICMP** or **ICMPv6**
- Select **OK** to confirm the configuration

Example

Example settings for an ICMP.service. In this case it has been set up for some special testing of ICMP packets.

Field	Value
Name	ICMP test #4
Comments	For testing of proprietary network scanner
Service Type	Firewall
Show in Service List	enabled
Category	Network Services

Field	Value
Protocol Options	
Protocol Type	ICMP
Type	7
Code	<left blank>

Creating a new ICMP service in the CLI

The following is the creation of the same service using the command line.

```
config firewall service custom
edit ICMP test4
    set comment "For testing of proprietary network scanner"
    set category Network Services
    set protocol ICMP
    set icmptype 7
end
end
```

IP

- For the **Protocol Type** field, choose **IP** from the drop down menu
- In the **Protocol Number** field enter the numeric value based on the information found in "[Protocol Number](#)" on [page 1](#)
- Select **OK** to confirm the configuration

Example

Example settings for an IP.service. In this case it has been set up to communicate via an old protocol called QNX

Field	Value
Name	QNX
Comments	For QNX communications to the Development Lab
Service Type	Firewall
Show in Service List	enabled
Category	Uncategorized
Protocol Options	
Protocol Type	IP
Protocol Number	106

Creating a new ICMP service in the CLI

The following is the creation of the same service using the command line.

```
config firewall service custom
edit ICMP test4
    set comment "For QNX communications to the Development Lab "
    set protocol IP
    set icmptype 106
end
end
```



In the CLI examples, the fields for **Show in Service List**, **Service Type** and in the example for IP, **Category** were not set because the values that they would have been set to were the default values and were already correctly set.

ALL/CONNECT/FTP/HTTP/SOCKS-TCP/SOCKS-UDP

These options are available only if the **Service Type** is set to **Explicit Proxy**.

8. For the **Protocol Type** field, choose one of the following from the drop down menu:
 - **ALL**
 - **CONNECT**
 - **FTP**
 - **HTTP**
 - **SOCKS-TCP**
 - **SOCKS-UDP**
9. For the **Address** field, choose IP Range or **FQDN** (Fully Qualified Domain Name) if there is to be a specific destination for the service. Depending on which type of address is selected, the field value needs to be filled with a FQDN string or an IP address in one of the 3 standard IPv4 address formats:
 - x.x.x.x - for a specific address
 - x.x.x.x/x - for a subnet
 - x.x.x.x-x.x.x.x - for a range of specific addresses
10. Configure the **Destination Port** by:
 - Enter the low end to the TCP port range in the field indicated by grayed out **Low**.
 - Enter the high end of the TCP port range in the field indicated by grayed out **High**. If there is only a single port in the range **High** can be left empty
 - Multiple ports or port ranges can be added by using the "+" at the beginning of the row
 - Rows can be removed by using the trash can symbol at the end of the row
11. If required, you can **Specify Source Ports** for the service by enabling the toggle switch.
 - The **Src Port** will match up with a **Destination Port**
 - **Src Ports** cannot be configured without there being a value for the **Destination Port**
 - The same rules for configuring the **Destination Ports** applies to the **Src Ports**
12. Select **OK** to confirm the configuration

Specific addresses in TCP/UDP/SCTP

In the TCP/UDP/SCTP services it is also possible to set the parameter for a specific IP or Fully Qualified Domain Name address. The IP/FQDN field refers to the destination address of the traffic, not the source. This means for example, that you can set up a custom service that will describe in a policy the TCP traffic over port 80 going to the web site example.com, but you cannot set up a service that describes the TCP traffic over port 80 that is coming from the computer with the address 192.168.29.59.

Service groups

Just like some of the other firewall components, services can also be bundled into groups for ease of administration.

Creating a service group

1. Go to **Policy & Objects > Services**.
2. Select **Create New**. A drop down menu is displayed. Select **Service Group**
3. Input a **Group Name** to describe the services being grouped
4. Input any additional information in the **Comments** field.
5. Choose a **Type** of group. The options are **Firewall** or **Explicit Proxy**.
6. Add to the list of **Members** from the drop down menu. Using the + sign beside the field will allow the addition of multiple services.
7. Press **OK**.

Example

Example of a New Service Group:

Field	Value
Group Name	Authentication Services
Comments	Services used in Authentication
Type	Firewall
Members	<ul style="list-style-type: none">• Kerberos• LDAP• LDAP_UDP• RADIUS

Firewall schedules

Firewall schedules control when policies are in effect. When you add a security policy on a FortiGate unit you need to set a schedule to determine the time frame in which that the policy will be functioning. While it is not set by default, the normal schedule would be always. This would mean that the policy that has been created is always function and always policing the traffic going through the FortiGate. The time component of the schedule is based on a 24 hour clock notation or military time as some people would say.

There are two types of schedules: One-time schedules and recurring schedules.

One-time schedule object

One-Time schedules are in effect only once for the period of time specified in the schedule. This can be useful for testing to limit how long a policy will be in effect in case it is not removed, or it can be used for isolated events such as a conference where you will only need a temporary infrastructure change for a few days.

The time frame for a One-time schedule is configured by using a start time which includes, Year | Month | Day | Hour | Minute and a Stop time which includes the same variables. So while the frequency of the schedule is only once it can last anywhere from 1 minute to multiple years.

Configuring a one-time schedule object in the GUI

1. Go to **Policy & Objects > Schedules**.
2. Select **Create New**. A drop down menu is displayed. Select **Schedule**.
3. From the **Type** options, choose **One-time**.
4. Input a **Name** for the schedule object.
5. If you wish to add a **Color** to the icon in the GUI, you can click on the **Change** link to choose 1 of 32 color options.
6. Choose a **Start Date**.
Selecting the field with the mouse will bring up a interactive calendar graphic that will allow the user to select the date. The date can also be typed in using the format YYYY/MM/DD.
7. Choose a **Start Time**.
The **Start Time** is composed of two fields, **Hour** and **Minute**. Think of setting the time for a digital clock in 24 hour mode. The **Hour** value can be an integer from 0 and 23. The **Minute** value can be from 0 to 59. 0 and 0 would be midnight at the start of the day and 23 and 59 would be one minute to midnight at the end of the day. The value can be entered by keyboard or by using the up and down arrows in the field to select the value.
6. Choose an **End Date**.
Configuration is the same as **Start Date**.
8. Choose a **Stop Time**.
Configuration is the same as **Start Time**.
9. Enable/Disable **Pre-expiration event log**.
This configures the system to create an event log 1 to 100 days before the **End Date** as a warning in case the schedule needs to be extended.
10. If the **Pre-expiration event log** is enabled, set the value for **Number of days before**.
11. Press **OK**.

Example: Firewall schedule - one-time

The company wants to change over their web site image to reference the new year. They have decided to take this opportunity to do some hardware upgrades as well. Their web site is business oriented so they have determined that over New Year's Eve there will be very limited traffic.

- They are going to need a maintenance window of 2 hours bracketing midnight on New Year's Eve.

Configuration in the GUI

1. Go to **Policy & Objects > Objects > Schedule**.
2. Select **Create New > Schedule**.
3. Fill out the fields with the following information:

Type	One-time
Name	NewYearsEve_Maintenance
Start Date	2014/12/31 <use the built in calendar>
End Date	2015/01/01 <use the built in calendar>
Start Time	Hour: 23, Minute: 0
Stop Time	Hour: 1Minute: 0
Pre-expiration event log	<disable>

4. Select **OK**.

To verify that the schedule was added correctly:

1. Go to **Policy & Objects > Objects > Schedule**.
2. Check that the schedule with the name you used has been added to the list of recurring schedules and that the listed settings are correct.

Configuration in the CLI

1. Enter the following CLI command:

```
config firewall schedule onetime
edit maintenance_window
set start 23:00 2012/12/31
set end 01:00 2013/01/01
next
end
```

To verify that the schedule was added correctly:

1. Enter the following CLI command:

```
config firewall schedule onetime
edit <the name of the schedule you wish to verify>
show full-configuration
```

Recurring schedule object

Recurring schedules are in effect repeatedly at specified times of specified days of the week. The Recurring schedule is based on a repeating cycle of the days of the week as opposed to every x days or days of the month. This means that you can configure the schedule to be in effect on Tuesday, Thursday, and Saturday but not every 2 days or on odd numbered days of the month.

If a recurring schedule has a stop time that is earlier than the start time, the schedule will take effect at the start time but end at the stop time on the next day. You can use this technique to create recurring schedules that run from one day to the next.

Configuring a recurring schedule object in the GUI

1. Go to **Policy & Objects > Schedules**.
2. Select **Create New**. A drop down menu is displayed. Select **Schedule**.
3. From the **Type** options, choose **Recurring**.
4. Input a **Name** for the schedule object.
5. If you wish to add a **Color** to the icon in the GUI, you can click on the **Change** link to choose 1 of 32 color options.
6. From the **Days** options, choose the day of the week that you would like this schedule to apply to. The schedule will be in effect on the days of the week that have a check mark in the checkbox to the left of the name of the weekday.
7. If the scheduled time is the whole day, leave the **All Day** toggle switch enabled. If the schedule is for specific times during the day, disable the **All Day** toggle switch.
8. If the All Day option is disabled, choose a **Start Time**.
The **Start Time** is composed of two fields, **Hour** and **Minute**. Think of setting the time for a digital clock in 24 hour mode. The **Hour** value can be an integer from 0 and 23. The **Minute** value can be from 0 to 59. 0 and 0 would be midnight at the start of the day and 23 and 59 would be one minute to midnight at the end of the day. The value can be entered by keyboard or by using the up and down arrows in the field to select the value.
7. Choose a **Stop Time**.
Configuration is the same as **Start Time**.
8. Press **OK**.



Because recurring schedules do not work with DENY policies, the strategy when designing a schedule should *not* be to determine when users cannot access a policy but to build the schedules around when it *is* possible to access the policy.

Example: Firewall schedule - recurring

The Company wants to allow the use of Facebook by employees, but only during none business hours and the lunch break.

- The business hours are 9:00 p.m. to 6:00 p.m.
- The Lunch break is 12:00 p.m. to 1:00 p.m.
- The plan is to create a schedule to cover the morning business hours and the afternoon business hours and block access to the Facebook web site during that time.

Configuration in the GUI

1. Go to **Policy & Objects > Objects > Schedule**.
2. Select **Create New > Schedule**.
3. Fill out the fields with the following information:

Type	Recurring
Name	Morning_Business_Hours
Days	Monday, Tuesday, Wednesday, Thursday, Friday
Start Time	Hour = 9, Minute = 0
Stop Time	Hour = 12, Minute = 0

4. Select **OK**.
5. Create a second new schedule.

Type	Recurring
Name	Morning_Business_Hours
Days	Monday, Tuesday, Wednesday, Thursday, Friday
Start Time	Hour = 13, Minute = 0
Stop Time	Hour = 18, Minute = 0

6. Select **OK**.

To verify that the schedule was added correctly:

1. Go to **Policy & Objects > Objects > Schedule**.
2. Check that the schedule with the name you used has been added to the list of recurring schedules and that the listed settings are correct.

Configuration in the CLI

1. Enter the following CLI command:

```
config firewall schedule recurring
edit Morning_Business_Hours
set day monday tuesday wednesday thursday friday
set start 09:00
set end 12:00
end
```

2. Enter the following CLI command:

```
config firewall schedule recurring
edit Afternoon_Business_Hours
set day monday tuesday wednesday thursday friday
set start 13:00
set end 18:00
end
```

To verify that the schedule was added correctly:

1. Enter the following CLI command:

```
config firewall schedule recurring
edit <the name of the schedule you wish to verify>
```



```
show full-configuration
```

Schedule groups

You can organize multiple firewall schedules into a schedule group to simplify your security policy list. The schedule parameter in the policy configuration does not allow for the entering of multiple schedules into a single policy so if you have a combination of time frames that you want to schedule the policy for then the best approach, rather than making multiple policies is to use a schedule group.

Creating a schedule group object

1. Go to **Policy & Objects > Schedules**.
2. Select **Create New**. A drop down menu is displayed. Select **Schedule Group**
3. Input a **Name** for the schedule object.
4. In the **Members** field, select the "+" to bring forth the panel for selecting entries.
5. Press **OK**.

Example

Your Internet policy allows employees to visit Social Media sites from company computers but not during what is considered working hours. The offices are open a few hours before working hours and the doors are not locked until a few hours after official closing so work hours are from 9 to 5 with a lunch break from Noon to 1:00 p.m.

Your approach is to block the traffic between 9 and noon and between 1:00 p.m. and 5:00 p.m. This means you will need two schedules for a single policy and the schedule group handles this for you. Schedule groups can contain both recurring and one-time schedules. Schedule groups cannot contain other schedule groups.

Schedule expiration

The schedule in a security policy enables certain aspects of network traffic to occur for a specific length of time. What it does not do however, is police that time. That is, the policy is active for a given time frame, and as long as the session is open, traffic can continue to flow.

For example, in an office environment, Skype use is allowed between noon and 1pm. During that hour, any Skype traffic continues. As long as that session is open, after the 1pm end time, the Skype conversations can continue, yet new sessions will be blocked. Ideally, the Skype session should close at 1pm.

Using a CLI command you can set the schedule to terminate all sessions when the end time of the schedule is reached. Within the config firewall command enter the command:

```
set schedule-timeout enable
```

By default, this option is set to disable.

A few further settings are needed to make this work.

```
config firewall policy
  edit ID
    set firewall-session-dirty check-new
  end

config system settings
  set firewall-session-dirty check-policy-option
end
```



The Policy window will indicate when a policy has become invalid due to its schedule parameters referring only to times in the past.

Firewall-session-dirty setting

The firewall-session-dirty setting has three options

<code>check-all</code>	CPU flushes all current sessions and re-evaluates them. [default]
<code>check-new</code>	CPU keeps existing sessions and applies policy changes to new sessions only. This reduces CPU load and the possibility of packet loss.
<code>check-policy-option</code>	Use the option selected in the firewall-session-dirty field of the firewall policy (check-all or check-new, as above, but per policy).

Secure web gateway, WAN optimization, web caching and WCCP

You can use FortiGate WAN optimization and web caching to improve performance and security of traffic passing between locations on your wide area network (WAN) or from the Internet to your web servers. You can also use the FortiGate unit as an explicit FTP and web proxy server. If your FortiGate unit supports web caching, you can also add web caching to any HTTP sessions including WAN optimization, explicit web proxy and other HTTP sessions.

the next sections of this document describes how FortiGate WAN optimization, web caching, explicit web proxy, explicit FTP proxy and WCCP work and also describes how to configure these features.

Before you begin

Before you begin to configure WAN optimization, Web caching, explicit proxies or WCCP, take a moment to note the following:

FortiGate models that support WAN optimization

WAN optimization is available on FortiGate models with internal storage that also support SSL acceleration. Internal storage includes high-capacity internal hard disks, AMC hard disk modules, FortiGate Storage Modules (FSMs) or over 4 Gbytes of internal flash storage. All of these storage locations can provide similar web caching and byte caching performance. If you add more than one storage location (for example, by creating multiple partitions on a storage device, by using more than one FSM, or by using an FSM and AMC hard disk in the same FortiGate unit) you can configure different storage locations for web caching and byte caching.

Distributing WAN optimization, explicit proxy, and web caching to multiple CPU cores

By default WAN optimization, explicit proxy and web caching is handled by half of the CPU cores in a FortiGate unit. For example, if your FortiGate unit has 4 CPU cores, by default two will be used for WAN optimization, explicit proxy and web caching. You can use the following command to change the number of CPU cores that are used.

```
config system global
    set wad-worker-count <number>
end
```

The value for <number> can be between 1 and the total number of CPU cores in your FortiGate unit. Adding more cores may enhance WAN optimization, explicit proxy and web caching performance and reduce the performance of other FortiGate systems.

Dispatching traffic to WAD worker based on source affinity

The `wad-worker` balancing algorithm supports a more balanced dispersal of traffic to the wad processes even, if the bulk of the traffic is coming from a small set of, or single source.

By default, dispatching traffic to WAD workers is based on source affinity. This may negatively affect performance when users have another explicit proxy in front of the FortiGate. Source affinity causes the FortiGate to process the traffic as if it originated from the single (or small set of) ip address of the outside proxy. This results in the use of one, or a small number, of WAD processes.

By disabling `wad-source-affinity` the traffic is balanced over all of the WAD processes. When the `wad-source-affinity` is disabled, the WAD dispatcher will not assign the traffic based on the source IP, but will assign the traffic to available workers in a round-robin fashion.



Handling the traffic by different WAD workers results in losing some of the benefits of using source affinity, as is explained by the warning message that appears when it is disabled:

"WARNING: Disabling this option results in some features to be unsupported. IP-based user authentication, disclaimer messages, security profile override, authentication cookies, MAPI scanning, and some video caches such as YouTube are not supported.

Do you want to continue? (y/n)"

CLI

```
config system global
    set wad-source-affinity {enable|disable}
end
```

Toggling disk usage for logging or wan-opt

Both logging and WAN Optimization use hard disk space to save data. In FortiOS, you cannot use the same hard disk for WAN Optimization and logging.

- If the FortiGate has one hard disk, then it can be used for either disk logging or WAN optimization, but not both. By default, the hard disk is used for disk logging.
- If the FortiGate has two hard disks, then one disk is always used for disk logging and the other disk is always used for WAN optimization.

On the FortiGate, go to **System > Advanced > Disk Settings** to switch between **Local Log** and **WAN Optimization**.

You can also change disk usage from the CLI using the following command:

```
configure system global
    set disk-usage {log | wanopt}
end
```



The Toggle Disk Usage feature is supported on all new "E" Series models, while support for "D" Series models may vary.

Please refer to the [Feature Platform Matrix](#) for more information.



Changing the disk setting formats the disk, erases current data stored on the disk and disables either disk logging or WAN Optimization.

You can configure WAN Optimization from the CLI or the GUI. To configure WAN Optimization from the GUI you must go to **System > Feature Visibility** and turn on WAN Optimization.



Remote logging (including logging to FortiAnalyzer and remote Syslog servers) is not affected by using the single local hard disk for WAN Optimization.

Enabling WAN optimization affects more than just disk logging

In addition to affecting WAN Optimization, the following table shows other features affected by the FortiGate disk configuration.

Features affected by Disk Usage as per the number of internal hard disks on the FortiGate

Feature	Logging Only (1 hard disk)	WAN Opt. Only (1 hard disk)	Logging & WAN Opt. (2 hard disks)
Logging	Supported	Not supported	Supported
Report/Historical FortiView	Supported	Not supported	Supported
Firewall Packet Capture (Policy Capture and Interface Capture)	Supported	Not supported	Supported
AV Quarantine	Supported	Not supported	Supported
IPS Packet Capture	Supported.	Not supported	Supported
DLP Archive	Supported	Not supported	Supported
Sandbox DB & Results	FortiSandbox database and results are also stored on disk, but will not be affected by this feature.		

Example topologies relevant to WAN optimization

FortiGate WAN optimization consists of a number of techniques that you can apply to improve the efficiency of communication across your WAN. These techniques include protocol optimization, byte caching, web caching, SSL offloading, and secure tunneling. Protocol optimization can improve the efficiency of traffic that uses the CIFS, FTP, HTTP, or MAPI protocol, as well as general TCP traffic. Byte caching caches files and other data on FortiGate units to reduce the amount of data transmitted across the WAN. Web caching stores web pages on FortiGate units to reduce latency and delays between the WAN and web servers. SSL offloading offloads SSL decryption and encryption from web servers onto FortiGate SSL acceleration hardware. Secure tunneling secures traffic as it crosses the WAN.

You can apply different combinations of these WAN optimization techniques to a single traffic stream depending on the traffic type. For example, you can apply byte caching and secure tunneling to any TCP traffic. For HTTP and HTTPS traffic, you can also apply protocol optimization and web caching.

You can configure a FortiGate unit to be an explicit web proxy server for both IPv4 and IPv6 traffic and an explicit FTP proxy server. Users on your internal network can browse the Internet through the explicit web proxy server or connect to FTP servers through the explicit FTP proxy server. You can also configure these proxies to protect access to web or FTP servers behind the FortiGate unit using a reverse proxy configuration.

Web caching can be applied to any HTTP or HTTPS traffic, this includes normal traffic accepted by a security policy, explicit web proxy traffic, and WAN optimization traffic.


You can also configure a FortiGate unit to operate as a Web Cache Communication Protocol (WCCP) client or server. WCCP provides the ability to offload web caching to one or more redundant web caching servers.

FortiGate units can also apply security profiles to traffic as part of a WAN optimization, explicit web proxy, explicit FTP proxy, web cache and WCCP configuration. Security policies that include any of these options can also include settings to apply all forms of security profiles supported by your FortiGate unit.

Basic WAN optimization topology

The basic FortiGate WAN optimization topology consists of two FortiGate units operating as WAN optimization peers intercepting and optimizing traffic crossing the WAN between the private networks.

Security device and WAN optimization topology



FortiGate units can be deployed as security devices that protect private networks connected to the WAN and also perform WAN optimization. In this configuration, the FortiGate units are configured as typical security devices for

the private networks and are also configured for WAN optimization. The WAN optimization configuration intercepts traffic to be optimized as it passes through the FortiGate unit and uses a WAN optimization tunnel with another FortiGate unit to optimize the traffic that crosses the WAN.


You can also deploy WAN optimization on single-purpose FortiGate units that only perform WAN optimization. In the out of path WAN optimization topology shown below, FortiGate units are located on the WAN outside of the private networks. You can also install the WAN optimization FortiGate units behind the security devices on the private networks.

The WAN optimization configuration is the same for FortiGate units deployed as security devices and for single-purpose WAN optimization FortiGate units. The only differences would result from the different network topologies.

Out-of-path WAN optimization topology


In an out-of-path topology, one or both of the FortiGate units configured for WAN optimization are not directly in the main data path. Instead, the out-of-path FortiGate unit is connected to a device on the data path, and the device is configured to redirect sessions to be optimized to the out-of-path FortiGate unit.

Single-purpose WAN optimization topology



The following out-of-path FortiGate units are configured for WAN optimization and connected directly to FortiGate units in the data path. The FortiGate units in the data path use a method such as policy routing to redirect traffic to be optimized to the out-of-path FortiGate units. The out-of-path FortiGate units establish a WAN optimization tunnel between each other and optimize the redirected traffic.

Out-of-path WAN optimization



One of the benefits of out-of-path WAN optimization is that out-of-path FortiGate units only perform WAN optimization and do not have to process other traffic. An in-path FortiGate unit configured for WAN optimization also has to process other non-optimized traffic on the data path.

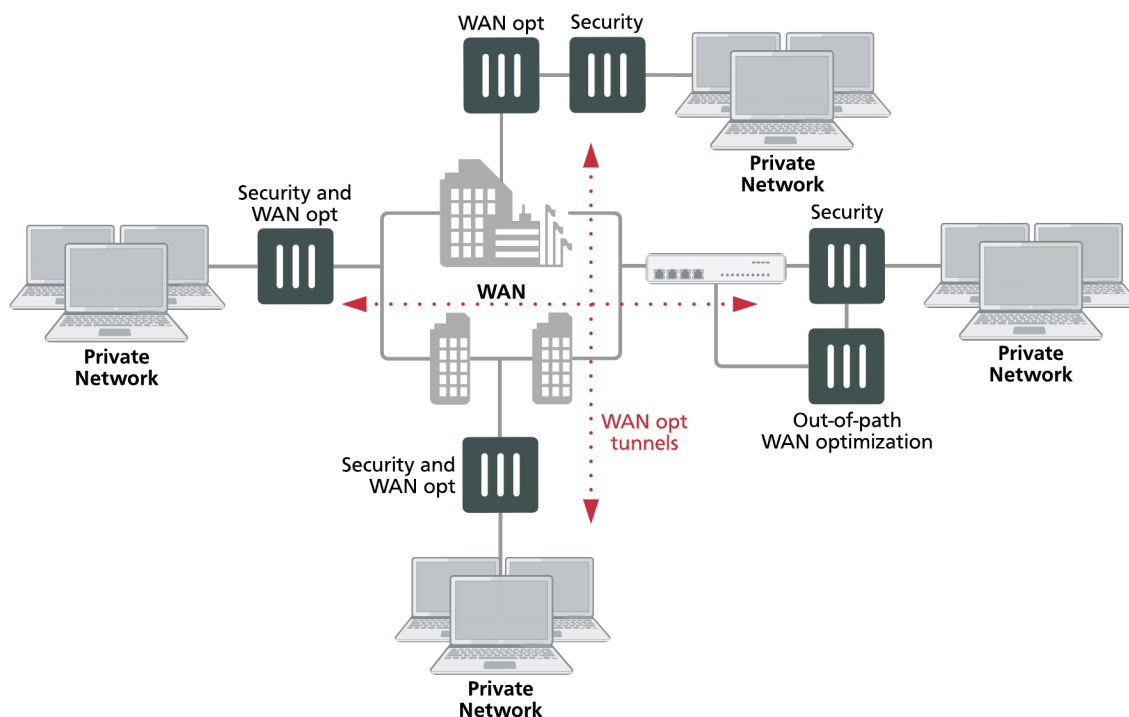
The out-of-path FortiGate units can operate in NAT/Route or transparent mode.

Other out-of-path topologies are also possible. For example, you can install the out-of-path FortiGate units on the private networks instead of on the WAN. Also, the out-of-path FortiGate units can have one connection to the network instead of two. In a one-arm configuration such as this, security policies and routing have to be configured to send the WAN optimization tunnel out the same interface as the one that received the traffic.

Topology for multiple networks

As shown in below, you can create multiple WAN optimization configurations between many private networks. Whenever WAN optimization occurs, it is always between two FortiGate units, but you can configure any FortiGate unit to perform WAN optimization with any of the other FortiGate units that are part of your WAN.

WAN optimization among multiple networks




You can also configure WAN optimization between FortiGate units with different roles on the WAN. FortiGate units configured as security devices and for WAN optimization can perform WAN optimization as if they are single-purpose FortiGate units just configured for WAN optimization.

WAN optimization with web caching

You can add web caching to a WAN optimization topology when users on a private network communicate with web servers located across the WAN on another private network.

WAN optimization with web caching topology



The topology above is the same as that shown in [WAN optimization with web caching on page 761](#) with the addition of web caching to the FortiGate unit in front of the private network that includes the web servers. You can also add web caching to the FortiGate unit that is protecting the private network. In a similar way, you can add web caching to any WAN Optimization topology.

Explicit web proxy topologies


You can configure a FortiGate unit to be an explicit web proxy server for Internet web browsing of IPv4 and IPv6 web traffic. To use the explicit web proxy, users must add the IP address of the FortiGate interface configured for the explicit web proxy to their web browser proxy configuration.

Explicit web proxy topology



If the FortiGate unit supports web caching, you can also add web caching to the security policy that accepts explicit web proxy sessions. The FortiGate unit then caches Internet web pages on a hard disk to improve web browsing performance.


Explicit web proxy with web caching topology



Explicit FTP proxy topologies


You can configure a FortiGate unit to be an explicit FTP proxy server for FTP users. To use the explicit web proxy, FTP users must connect to and authenticate with the explicit FTP proxy before connecting to an FTP server.

Explicit FTP proxy topology



You can also configure reverse explicit FTP proxy. In this configuration, users on the Internet connect to the explicit web proxy before connecting to an FTP server installed behind a FortiGate unit.

Reverse explicit FTP proxy topology




Web caching topologies

FortiGate web caching can be added to any security policy and any HTTP or HTTPS traffic accepted by that security policy can be cached on the FortiGate unit hard disk. This includes WAN optimization and explicit web proxy traffic. The network topologies for these scenarios are very similar. They involved a FortiGate unit installed between users and web servers with web caching enabled.


A typical web-caching topology includes one FortiGate unit that acts as a web cache server. Web caching is enabled in a security policy and the FortiGate unit intercepts web page requests accepted by the security policy, requests web pages from the web servers, caches the web page contents, and returns the web page contents to the users. When the FortiGate unit intercepts subsequent requests for cached web pages, the FortiGate unit contacts the destination web server just to check for changes.

Web caching topology



You can also configure reverse proxy web-caching. In this configuration, users on the Internet browse to a web server installed behind a FortiGate unit. The FortiGate unit intercepts the web traffic (HTTP and HTTPS) and caches pages from the web server. Reverse proxy web caching on the FortiGate unit reduces the number of requests that the web server must handle, leaving it free to process new requests that it has not serviced before.

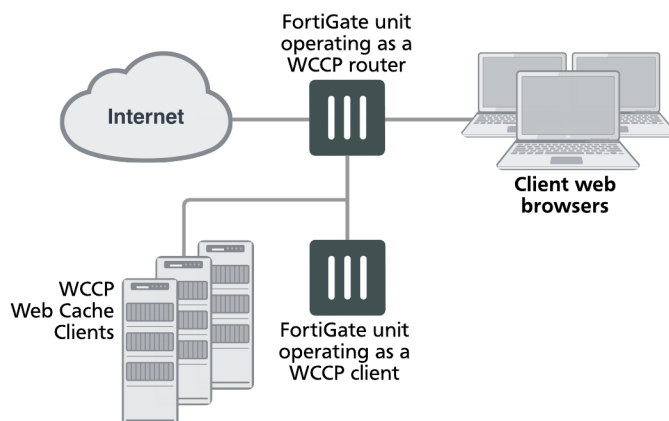
Reverse proxy web caching topology



WCCP topologies

You can operate a FortiGate unit as a Web Cache Communication Protocol (WCCP) router or cache engine. As a router, the FortiGate unit intercepts web browsing requests from client web browsers and forwards them to a WCCP cache engine. The cache engine returns the required cached content to the client web browser. If the cache server does not have the required content it accesses the content, caches it and returns the content to the client web browser.

WCCP topology



FortiGate units can also operate as WCCP cache servers, communicating with WCCP routers, caching web content and providing it to client web browsers as required.

WCCP is transparent to client web browsers. The web browsers do not have to be configured to use a web proxy.

Inside FortiOS: WAN optimization

Enterprises deploying FortiOS can leverage WAN optimization to provide fast and secure application responses between locations on a Wide Area Network (WAN). The web caching component of FortiOS WAN optimization extends this protection and performance boost to cloud services.

Centralize without compromising your WAN performance

Many multi-location enterprise environments reduce costs and consolidate resources by centralizing applications or providing applications in the cloud. Efficient and high-speed communication between applications and their users is critical. Remote sites don't always have access to high bandwidth, but users at all sites expect consistent network performance. Minimizing user impact and improving performance is especially vital when applications designed for local area networks (LANs) are on the cloud.

Even applications that work fine on a local LAN, such as Windows File Sharing (CIFS), email exchange (MAPI), and many others, suffer from bandwidth limitations and latency issues when accessed over a WAN. This results in a loss of productivity and a perceived need for expensive network upgrades. FortiOS's WAN Optimization provides an inexpensive and easy way to deploy a solution to this problem.



FortiOS is commonly deployed in central offices, satellite offices, and in the cloud to provide secure communications across a WAN using IPsec or SSL VPN. This installed infrastructure can be leveraged to add more value by using WAN Optimization to accelerate WAN traffic and web caching to accelerate cloud services.

FortiOS WAN optimization

FortiOS includes license-free WAN optimization on most current FortiGate devices. WAN optimization is a comprehensive solution that maximizes your WAN performance and provides intelligent bandwidth management and unmatched consolidated security performance. WAN optimization reduces your network overhead and removes unnecessary traffic for a better overall performance experience. Efficient use of bandwidth and better application performance will remove the need for costly WAN link upgrades between data centers and other expensive solutions for your network traffic growth.

Protocol optimization

Protocol optimization is effective for applications designed for the LAN that do not function well on low bandwidth high latency networks. FortiOS protocol optimization improves the efficiency of CIFS, FTP, HTTP, MAPI, and general TCP sessions.

Regular bandwidth usage**Improved bandwidth usage with FortiGate protocol optimization**

For example, CIFS, which is a fairly “chatty” protocol, requires many background transactions to successfully transfer a single file. When transferring the file, CIFS sends small chunks of data and waits sequentially for each chunk’s arrival and acknowledgment before sending the next. This large amount of request/acknowledgement traffic can delay transfers. FortiOS CIFS WAN Optimization removes this chattiness and gets on with the job of transferring the file.

TCP protocol optimization uses techniques such as SACK support, window scaling and window size adjustment, and connection pooling to remove common WAN TCP bottlenecks.

Web caching

In an enterprise environment, multiple users will often want to get the same content (for example, a sales spreadsheet, a corporate presentation or a PDF from a cloud service, or a software update). With FortiOS Web caching, content from the cloud, from the web or from other sites on the WAN is download once and cached on the local FortiGate device. When other users access the same content they download it from the cache. The result is less bandwidth use and reduced latency for the file requester.

FortiOS web caching also recognizes requests for Windows or MS-Office updates and downloads the new update file in the background. Once downloaded to the cache, the new update file is available to all users and all subsequent requests for this update are rapidly downloaded from the cache.

Byte caching

Byte caching improves caching by accelerating the transfer of similar, but not identical content. Byte caching accelerates multiple downloads of different email messages with the same corporate disclaimer by downloading the disclaimer over the WAN once and then downloading all subsequent disclaimers from a local FortiGate unit. Byte caching reduces the amount of data crossing the WAN when multiple different emails with the same or similar attachments or different versions of an attachment are downloaded from a corporate email server to different locations over the WAN.

Dynamic data chunking

Dynamic data chunking detects and optimizes persistent data chunks in changed files or in data embedded in traffic that uses an unknown protocol. For example, dynamic chunking can cache data in Lotus notes traffic and make the data chunks available for email and other protocols.

Data deduplication

Byte caching breaks large units of application data, like an email attachment or a file download, into manageable small chunks of data. Each chunk of data is labeled with a hash, and chunks with their respective hashes are stored in a database on the local FortiGate unit. When a remote user request a file, the WAN Optimization sends the hashes, rather than the actual data. The FortiGate unit at the other end of the WAN tunnel reassembles the data from its own hash database, only downloading chunks that it is missing. Deduplication, or the process of eliminating duplicate data, will reduce space consumption. In addition to reducing the amount of data downloaded across the WAN, byte caching is not application specific and assists by accelerating all of the protocols supported by WAN Optimization.

Server monitoring and management

The health and performance of real servers can be monitored from the FortiGate GUI. Virtual servers and their assigned real servers can be monitored for health status, if there have been any monitor events, number of active sessions, round trip time and number of bytes processed. Should a server become problematic and require administration, it can be gracefully removed from the Real Server pool to enable disruption free maintenance. When a removed real server is able to operate it can gracefully be added back to the virtual server.

SSL acceleration

SSL is used by many organizations to keep WAN communications private. WAN Optimization boosts SSL acceleration properties of FortiGate FortiASIC hardware by accelerating SSL traffic across the WAN. The FortiGate unit handles SSL encryption/decryption for corporate servers providing SSL encrypted connections over the WAN.

VPN replacement

FortiOS WAN optimization supports secure SSL-encrypted tunnels between FortiGate units on the WAN. Employing secure WAN Optimization tunnels can replace IPsec VPNs between sites. The result is a single, relatively simple configuration that supports optimization and privacy of communication across the WAN and uses FortiGate SSL acceleration to provide high performance.

Road warriors and home workers

The drive to give employees greater flexibility and reduce operational costs has led to more remote workers, both at home and on the road. Whether accessing the office from a hotel, public wireless hotspot, or home, the problem is the same: low bandwidth and high latency harming application performance. WAN Optimization is integrated into FortiClient, which can be installed on PCs and wireless devices to optimize communication between remote workers and their offices.

Reduce your...

- **Capital outlay:** Organizations only need to purchase a single device per location.
- **Licensing costs:** WAN Optimization is included with FortiOS. Additional licenses are not needed.
- **Network complexity:** Small offices that may not have the space or power connections for multiple devices do not need to worry: no additional devices are required.

WAN optimization concepts

Client/server architecture

Traffic across a WAN typically consists of clients on a client network communicating across a WAN with a remote server network. The clients do this by starting communication sessions from the client network to the server network. These communication sessions can be open text over the WAN or they can be encrypted by SSL VPN or IPsec VPN.

To optimize these sessions, you can add **WAN optimization security policies** to the **client-side FortiGate unit** to accept sessions from the client network that are destined for the server network. The client-side FortiGate unit is located between the client network and the WAN. WAN optimization security policies include **WAN optimization profiles** that control how the traffic is optimized.


The client-side FortiGate unit must also include the IP address of the **server-side FortiGate unit** in its WAN optimization **peer** configuration. The server-side FortiGate unit is located between the server network and the WAN. The peer configuration allows the client-side FortiGate unit to find the server-side FortiGate unit and attempt to establish a WAN optimization **tunnel** with it.

For the server-side FortiGate unit you must add a security policy with **wanopt** as the **Incoming Interface**. This security policy allows the FortiGate unit to accept WAN optimization sessions from the client-side FortiGate unit. For the server-side FortiGate unit to accept a WAN optimization connection it must have the client-side FortiGate unit in its WAN optimization peer configuration.



WAN optimization profiles are only added to the client-side WAN optimization security policy. The server-side FortiGate unit employs the WAN optimization settings set in the WAN optimization profile on the client-side FortiGate unit.

Client/server architecture



When both peers are identified the FortiGate units attempt to establish a WAN optimization **tunnel** between them. WAN optimization tunnels use port 7810. All optimized data flowing across the WAN between the client-side and server-side FortiGate units use this tunnel. WAN optimization tunnels can be encrypted using SSL encryption to keep the data in the tunnel secure.

Any traffic can be sent through a WAN optimization tunnel. This includes SSL and IPsec VPN traffic. However, instead of configuring SSL or IPsec VPN for this communication you can add SSL encryption using the WAN optimization tunnel.

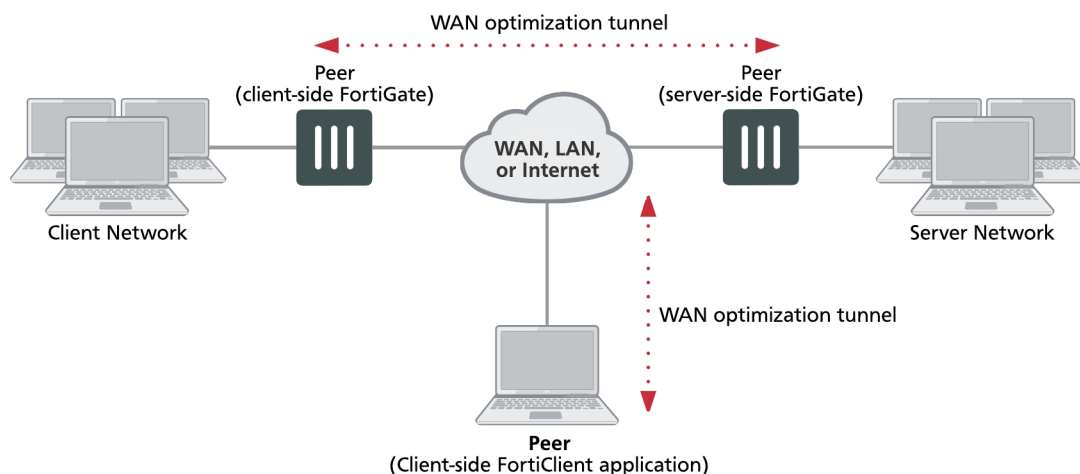
In addition to basic identification by peer host ID and IP address you can configure WAN optimization **authentication** using certificates and pre-shared keys to improve security. You can also configure FortiGate units involved in WAN optimization to accept connections from any identified peer or restrict connections to specific peers.

The FortiClient application can act in the same manner as a client-side FortiGate unit to optimize traffic between a computer running FortiClient and a FortiGate unit.

WAN optimization peers

The client-side and server-side FortiGate units are called WAN optimization peers because all of the FortiGate units in a WAN optimization network have the same peer relationship with each other. The client and server roles just relate to how a session is started. Any FortiGate unit configured for WAN optimization can be a client-side and a server-side FortiGate unit at the same time, depending on the direction of the traffic. Client-side FortiGate units initiate WAN optimization sessions and server-side FortiGate units respond to the session requests. Any FortiGate unit can simultaneously be a client-side FortiGate unit for some sessions and a server-side FortiGate unit for others.

WAN optimization peer and tunnel architecture



To identify all of the WAN optimization peers that a FortiGate unit can perform WAN optimization with, you add host IDs and IP addresses of all of the peers to the FortiGate unit configuration. The peer IP address is actually the IP address of the peer unit interface that communicates with the FortiGate unit.

Protocol optimization

Protocol optimization techniques optimize bandwidth use across the WAN. These techniques can improve the efficiency of communication across the WAN optimization tunnel by reducing the amount of traffic required by communication protocols. You can apply protocol optimization to Common Internet File System (CIFS), FTP, HTTP, MAPI, and general TCP sessions. You can apply general TCP optimization to MAPI sessions.

For example, CIFS provides file access, record locking, read/write privileges, change notification, server name resolution, request batching, and server authentication. CIFS is a fairly “chatty” protocol, requiring many background transactions to successfully transfer a single file. This is usually not a problem across a LAN. However, across a WAN, latency and bandwidth reduction can slow down CIFS performance.

When you select the CIFS protocol in a WAN optimization profile, the FortiGate units at both ends of the WAN optimization tunnel use a number of techniques to reduce the number of background transactions that occur over the WAN for CIFS traffic.

If a policy accepts a range of different types of traffic, you can set **Protocol** to **TCP** to apply general optimization techniques to TCP traffic. However, applying this TCP optimization is not as effective as applying more protocol-specific optimization to specific types of traffic. TCP protocol optimization uses techniques such as TCP SACK support, TCP window scaling and window size adjustment, and TCP connection pooling to remove TCP bottlenecks.

Protocol optimization and MAPI

By default the MAPI service uses port number 135 for RPC port mapping and may use random ports for MAPI messages. The random ports are negotiated through sessions using port 135. The FortiOS DCE-RPC session helper learns these ports and opens pinholes for the messages. WAN optimization is also aware of these ports and attempts to apply protocol optimization to MAPI messages that use them. However, to configure protocol optimization for MAPI you should set the WAN optimization profile to a single port number (usually port 135). Specifying a range of ports may reduce performance.

Byte caching

Byte caching breaks large units of application data (for example, a file being downloaded from a web page) into small chunks of data, labeling each chunk of data with a hash of the chunk and storing those chunks and their hashes in a database. The database is stored on a WAN optimization storage device. Then, instead of sending the actual data over the WAN tunnel, the FortiGate unit sends the hashes. The FortiGate unit at the other end of the tunnel receives the hashes and compares them with the hashes in its local byte caching database. If any hashes match, that data does not have to be transmitted over the WAN optimization tunnel. The data for any hashes that does not match is transferred over the tunnel and added to that byte caching database. Then the unit of application data (the file being downloaded) is reassembled and sent to its destination.

The stored byte caches are not application specific. Byte caches from a file in an email can be used to optimize downloading that same file or a similar file from a web page.

The result is less data transmitted over the WAN. Initially, byte caching may reduce performance until a large enough byte caching database is built up.

To enable byte caching, you select **Byte Caching** in a WAN optimization profile.

Byte caching cannot determine whether or not a file is compressed (for example a zip file), and caches compressed and non-compressed versions of the same file separately.

Dynamic data chunking for byte caching

Dynamic data chunking can improve byte caching by improving detection of data chunks that are already cached in changed files or in data embedded in traffic using an unknown protocol. Dynamic data chunking is available for HTTP, CIFS and FTP.

Use the following command to enable dynamic data chunking for HTTP in the default WAN optimization profile.

```
config wanopt profile
  edit default
    config http
      set prefer-chunking dynamic
    end
```

By default dynamic data chunking is disabled and `prefer-chunking` is set to `fix`.

WAN optimization transparent mode

WAN optimization is transparent to users. This means that with WAN optimization in place, clients connect to servers in the same way as they would without WAN optimization. However, servers receiving packets after WAN optimization “see” different source addresses depending on whether or not transparent mode is selected for WAN optimization. If transparent mode is selected, WAN optimization keeps the original source address of the packets, so servers appear to receive traffic directly from clients. Routing on the server network should be configured to route traffic with client source IP addresses from the server-side FortiGate unit to the server and back to the server-side FortiGate unit.



Some protocols, for example CIFS, may not function as expected if transparent mode is **not** selected. In most cases, for CIFS WAN optimization you should select transparent mode and make sure the server network can route traffic as described to support transparent mode.

If transparent mode is not selected, the source address of the packets received by servers is changed to the address of the server-side FortiGate unit interface that sends the packets to the servers. So servers appear to receive packets from the server-side FortiGate unit. Routing on the server network is simpler in this case because client addresses are not involved. All traffic appears to come from the server-side FortiGate unit and not from individual clients.



Do not confuse WAN optimization transparent mode with FortiGate transparent mode. WAN optimization transparent mode is similar to source NAT. FortiGate's transparent mode is a system setting that controls how the FortiGate unit (or a VDOM) processes traffic.

Configuring transparent mode

You can configure transparent mode by selecting **Transparent** in a WAN Optimization profile. The profile is added to an active WAN Optimization policy.

When you configure a passive WAN Optimization policy you can accept the active policy transparent setting or you can override the active policy transparent setting. From the GUI you can do this by setting the **Passive Option** as follows:

- **default** use the transparent setting in the WAN Optimization profile added to the active policy (client-side configuration).
- **transparent** impose transparent mode (override the active policy transparent mode setting). Packets exiting the FortiGate keep their original source addresses.
- **non-transparent** impose non-transparent mode (override the active policy transparent mode setting). Packets exiting the FortiGate have their source address changed to the address of the server-side FortiGate unit interface that sends the packets to the servers.

From the CLI you can use the following command:

```
config firewall policy
    set wanopt-passive-opt {default | transparent | non-transparent}
end
```

Operating modes and VDOMs

To use WAN optimization, the FortiGate units can operate in either NAT/Route or transparent mode. The client-side and server-side FortiGate units do not have to be operating in the same mode.

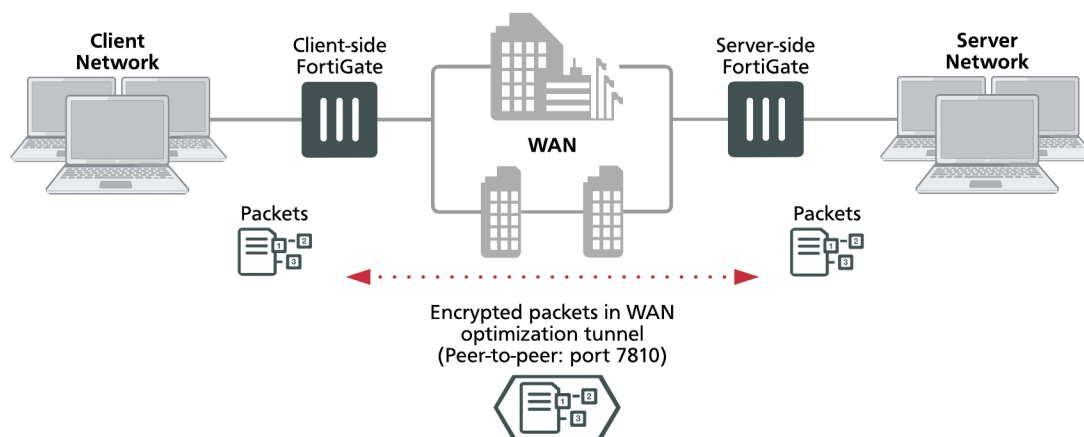
As well, the FortiGate units can be configured for multiple virtual domain (VDOM) operation. You configure WAN optimization for each VDOM and configure one or both of the units to operate with multiple VDOMs enabled.

If a FortiGate unit or VDOM is operating in transparent mode with WAN optimization enabled, WAN optimization uses the management IP address as the peer IP address of the FortiGate unit instead of the address of an interface.

WAN optimization tunnels

All optimized traffic passes between the FortiGate units over a WAN optimization tunnel. Traffic in the tunnel can be sent in plain text or encrypted using AES-128bit-CBC SSL.

WAN optimization tunnels



Both plain text and the encrypted tunnels use TCP destination port 7810.

Before a tunnel can be started, the peers must be configured to authenticate with each other. Then, the client-side peer attempts to start a WAN optimization tunnel with the server-side peer. Once the peers authenticate with each other, they bring up the tunnel and WAN optimization communication over the tunnel starts. After a tunnel has been established, multiple WAN optimization sessions can start and stop between peers without restarting the tunnel.

Tunnel sharing

You can use the `tunnel-sharing` WAN optimization profile CLI keyword to configure tunnel sharing for WAN optimization rules. Tunnel sharing means multiple WAN optimization sessions share the same tunnel. Tunnel sharing can improve performance by reducing the number of WAN optimization tunnels between FortiGate units. Having fewer tunnels means less data to manage. Also, tunnel setup requires more than one exchange of information between the ends of the tunnel. Once the tunnel is set up, each new session that shares the tunnel avoids tunnel setup delays.

Tunnel sharing also uses bandwidth more efficiently by reducing the chances that small packets will be sent down the tunnel. Processing small packets reduces network throughput, so reducing the number of small packets

improves performance. A shared tunnel can combine all the data from the sessions being processed by the tunnel and send the data together. For example, suppose a FortiGate unit is processing five WAN optimization sessions and each session has 100 bytes to send. If these sessions use a shared tunnel, WAN optimization combines the packets from all five sessions into one 500-byte packet. If each session uses its own private tunnel, five 100-byte packets will be sent instead. Each packet also requires a TCP ACK reply. The combined packet in the shared tunnel requires one TCP ACK packet. The separate packets in the private tunnels require five.

Use the following command to configure tunnel sharing for HTTP traffic in a WAN optimization profile.

```
config wanopt profile
  edit default
    config http
      set tunnel-sharing {express-shared | private | shared}
    end
```

Tunnel sharing is not always recommended and may not always be the best practice. Aggressive and non-aggressive protocols should not share the same tunnel. An aggressive protocol can be defined as a protocol that is able to get more bandwidth than a non-aggressive protocol. (The aggressive protocols can “starve” the non-aggressive protocols.) HTTP and FTP are considered aggressive protocols. If aggressive and non-aggressive protocols share the same tunnel, the aggressive protocols may take all of the available bandwidth. As a result, the performance of less aggressive protocols could be reduced. To avoid this problem, rules for HTTP and FTP traffic should have their own tunnel. To do this, set `tunnel-sharing` to `private` for WAN optimization rules that accept HTTP or FTP traffic.

It is also useful to set `tunnel-sharing` to `express-shared` for applications, such as Telnet, that are very interactive but not aggressive. Express sharing optimizes tunnel sharing for Telnet and other interactive applications where latency or delays would seriously affect the user’s experience with the protocol.

Set `tunnel-sharing` to `shared` for applications that are not aggressive and are not sensitive to latency or delays. WAN optimization rules set to `sharing` and `express-shared` can share the same tunnel.

WAN optimization and user and device identity policies, load balancing and traffic shaping

Please note the following about WAN optimization and firewall policies:

- WAN optimization is not compatible with firewall load balancing.
- WAN optimization is compatible with source and destination NAT options in firewall policies (including firewall virtual IPs). If a virtual IP is added to a policy the traffic that exits the WAN optimization tunnel has its destination address changed to the virtual IPs mapped to IP address and port.
- WAN optimization is compatible with user identity-based and device identity security policies. If a session is allowed after authentication or device identification the session can be optimized.

Traffic shaping

Traffic shaping works for WAN optimization traffic that is not in a WAN optimization tunnel. So traffic accepted by a WAN optimization security policy on a client-side FortiGate unit can be shaped on ingress. However, when the traffic enters the WAN optimization tunnel, traffic shaping is not applied.

In manual mode:

- Traffic shaping works as expected on the client-side FortiGate unit.
- Traffic shaping cannot be applied to traffic on the server-side FortiGate unit.

In active-passive mode:

- Traffic shaping works as expected on the client-side FortiGate unit.
- If transparent mode is enabled in the WAN optimization profile, traffic shaping also works as expected on the server-side FortiGate unit.
- If transparent mode is not enabled, traffic shaping works partially on the server-side FortiGate unit.

WAN optimization and HA

You can configure WAN optimization on a FortiGate HA cluster. The recommended best practice HA configuration for WAN optimization is active-passive mode. When the cluster is operating, all WAN optimization sessions are processed by the primary unit only. Even if the cluster is operating in active-active mode, HA does not load-balance WAN optimization sessions.

You can also form a WAN optimization tunnel between a cluster and a standalone FortiGate unit or between two clusters.

In a cluster, only the primary unit stores the byte cache database. This database is not synchronized to the subordinate units. So, after a failover, the new primary unit must rebuild its byte cache. Rebuilding the byte cache can happen relatively quickly because the new primary unit gets byte cache data from the other FortiGate unit that it is participating with in WAN optimization tunnels.

WAN optimization, web caching and memory usage

To accelerate and optimize disk access and to provide better throughput and less latency FortiOS WAN optimization uses provisioned memory to reduce disk I/O and increase disk I/O efficiency. In addition, WAN optimization requires a small amount of additional memory per session for comprehensive flow control logic and efficient traffic forwarding.

When WAN optimization is enabled you will see a reduction in available memory. The reduction increases when more WAN optimization sessions are being processed. If you are thinking of enabling WAN optimization on an operating FortiGate unit, make sure its memory usage is not maxed out during high traffic periods.

In addition to using the system dashboard to see the current memory usage you can use the `get test wad 2` command to see how much memory is currently being used by WAN optimization. See "get test {wad | wccpd} <test_level>" for more information.

WAN optimization configuration

This chapter describes FortiGate WAN optimization client server architecture and other concepts you need to understand to be able to configure FortiGate WAN optimization.

Manual (peer-to-peer) and active-passive WAN optimization

You can create **manual** (peer-to-peer) and **active-passive** WAN optimization configurations.



In reality, because WAN optimization traffic can only be processed by one CPU core, it is not recommended to increase the number of manual mode peers on the FortiGate unit per VDOM.

Note that the maximum number of manual peers are restricted to 256 per VDOM. However, in Active-Passive configurations, there is no hard-limit to the maximum number of manual peers per VDOM.

Manual (peer to peer) configurations

Manual configurations allow for WAN optimization between one client-side FortiGate unit and one server-side FortiGate unit. To create a manual configuration you add a **manual mode** WAN optimization security policy to the client-side FortiGate unit. The manual mode policy includes the peer ID of a server-side FortiGate unit.

In a manual mode configuration, the client-side peer can only connect to the named server-side peer. When the client-side peer initiates a tunnel with the server-side peer, the packets that initiate the tunnel include extra information so that the server-side peer can determine that it is a peer-to-peer tunnel request. This extra information is required because the server-side peer does not require a WAN optimization policy; however, you need to add the client peer host ID and IP address to the server-side FortiGate unit peer list.

In addition, from the server-side FortiGate unit CLI you must add an Explicit Proxy security policy with `proxy` set to `wanopt` and the destination interface and network set to the network containing the servers that clients connect to over the WAN optimization tunnel. WAN optimization tunnel requests are accepted by the explicit proxy policy and if the client-side peer is in the server side peer's address list the traffic is forwarded to the servers on the destination network.

Manual mode client-side policy

You must configure manual mode client-side policies from the CLI. From the GUI a manual mode policy has WAN Optimization turned on and includes the following text beside the *WAN optimization* field: *Manual (Profile: <profile-name>. Peer: <peer-name>.*

Add a manual mode policy to the client-side FortiGate unit from the CLI. The policy enables WAN optimization, sets `wanopt-detection` to `off`, and uses the `wanopt-peer` option to specify the server-side peer. The following example uses the default WAN optimization profile.

```
config firewall policy
  edit 2
    set srcintf internal
    set dstintf wan1
    set srcaddr client-subnet
    set dstaddr server-subnet
```

```

        set action accept
        set schedule always
        set service ALL
        set wanopt enable
        set wanopt-detection off
        set wanopt-profile default
        set wanopt-peer server
    next
end

```

Manual mode server-side explicit proxy policy

The server-side explicit proxy policy allows connections from the WAN optimization tunnel to the server network by setting the proxy type to `wanopt`. You must add policies that set `proxy` to `wanopt` from the CLI and these policies do not appear on the GUI. The policy should look like the following:

```

configure firewall proxy-policy
edit 3
    set proxy wanopt
    set dstintf internal
    set srcaddr all
    set dstaddr server-subnet
    set action accept
    set schedule always
    set service ALL
next
end

```

Active-passive configurations

Active-passive WAN optimization requires an **active** WAN optimization policy on the client-side FortiGate unit and a **passive** WAN optimization policy on the server-side FortiGate unit. The server-side FortiGate unit also requires an explicit proxy policy with `proxy` set to `wanopt`.

You can use the passive policy to control WAN optimization address translation by specifying **transparent mode** or non-transparent mode. See [Manual \(peer-to-peer\) and active-passive WAN optimization on page 776](#). You can also use the passive policy to apply security profiles, web caching, and other FortiGate features at the server-side FortiGate unit. For example, if a server-side FortiGate unit is protecting a web server, the passive policy could enable web caching.

A single passive policy can accept tunnel requests from multiple FortiGate units as long as the server-side FortiGate unit includes their peer IDs and all of the client-side FortiGate units include the server-side peer ID.

Active client-side policy

Add an active policy to the client-side FortiGate unit by turning on **WAN Optimization** and selecting **active**. Then select a WAN optimization **Profile**. From the CLI the policy could look like the following:

```

config firewall policy
edit 2
    set srcintf internal
    set dstintf wan1
    set srcaddr client-subnet
    set dstaddr server-subnet
    set action accept
    set schedule always
    set service ALL

```



```
        set wanopt enable
        set wanopt-detection active
        set wanopt-profile default
    next
end
```

Server-side tunnel policy

The server-side requires an explicit proxy policy that sets the `proxy` to `wanopt`. You must add this policy from the CLI and policies with `proxy` set to `wanopt` do not appear on the GUI. From the CLI the policy could look like the following:

```
configure firewall proxy-policy
edit 3
    set proxy wanopt
    set dstintf internal
    set srcaddr all
    set dstaddr server-subnet
    set action accept
    set schedule always
    set service ALL
next
end
```

Server-side passive policy

Add a passive policy to the server-side FortiGate unit by selecting **Enable WAN Optimization** and selecting **passive**. Then set the **Passive Option** to **transparent**. From the CLI the policy could look like the following:

```
config firewall policy
edit 2
    set srcintf "wan1"
    set dstintf "internal"
    set srcaddr "all"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY"
    set wanopt enable
    set wanopt-detection passive
    set wanopt-passive-opt transparent
next
```

WAN optimization profiles

Use WAN optimization profiles to apply WAN optimization techniques to traffic to be optimized. In a WAN optimization profile you can select the protocols to be optimized and for each protocol you can enable SSL offloading (if supported), secure tunneling, byte caching and set the port or port range the protocol uses. You can also enable transparent mode and optionally select an authentication group. You can edit the default WAN optimization profile or create new ones.

To configure a WAN optimization profile go to **WAN Opt. & Cache > Profiles** and edit a profile or create a new one.

Configuring a WAN optimization profile

Edit WAN Optimization Profile

default

Name

default

Comments

default WANopt profile

22/255

☒ Transparent Mode

☒ Authentication Group

Auth-Grp

Protocol	SSL Offloading	Secure Tunneling	Byte Caching	Port
<input checked="" type="checkbox"/> CIFS		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	445
<input checked="" type="checkbox"/> FTP		<input checked="" type="checkbox"/>	<input checked="" type="checkbox"/>	21
<input checked="" type="checkbox"/> HTTP	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input checked="" type="checkbox"/>	80
<input checked="" type="checkbox"/> MAPI		<input type="checkbox"/>	<input checked="" type="checkbox"/>	135
<input checked="" type="checkbox"/> TCP	<input checked="" type="checkbox"/>	<input type="checkbox"/>	<input type="checkbox"/>	1-65535

From the CLI you can use the following command to configure a WAN optimization profile to optimize HTTP traffic.

```
config wanopt profile
  edit new-profile
    config http
      set status enable
  end
```

Transparent Mode

Servers receiving packets after WAN optimization “see” different source addresses depending on whether or not you select **Transparent Mode**.

For more information, see [WAN optimization profiles on page 778](#).

Authentication Group

Select this option and select an authentication group so that the client and server-side FortiGate units must authenticate with each other before starting the WAN optimization tunnel. You must also select an authentication group if you select **Secure Tunneling** for any protocol.

You must add identical authentication groups to both of the FortiGate units that will participate in the WAN optimization tunnel. For more information, see [Configuring authentication groups on page 1](#).

Protocol

Select CIFS, FTP, HTTP or MAPI to apply protocol optimization for the selected protocols. See [WAN optimization profiles on page 778](#).

Select TCP if the WAN optimization tunnel accepts sessions that use more than one protocol or that do not use the CIFS, FTP, HTTP, or MAPI protocol.

SSL Offloading	<p>Select to apply SSL offloading for HTTPS or other SSL traffic. You can use SSL offloading to offload SSL encryption and decryption from one or more HTTP servers to the FortiGate unit. If you enable this option, you must configure the security policy to accept SSL-encrypted traffic.</p> <p>If you enable SSL offloading, you must also use the CLI command <code>config firewall ssl-server</code> to add an SSL server for each HTTP server that you want to offload SSL encryption/decryption for. For more information, see Turning on web caching for HTTPS traffic on page 1.</p>
Secure Tunneling	<p>The WAN optimization tunnel is encrypted using SSL encryption. You must also add an authentication group to the profile. For more information, see Secure tunneling on page 1.</p>
Byte Caching	<p>Select to apply WAN optimization byte caching to the sessions accepted by this rule. For more information, see "Byte caching".</p>
Port	<p>Enter a single port number or port number range. Only packets whose destination port number matches this port number or port number range will be optimized.</p>

Processing non-HTTP sessions accepted by a WAN optimization profile with HTTP optimization

From the CLI, you can use the following command to configure how to process non-HTTP sessions when a rule configured to accept and optimize HTTP traffic accepts a non-HTTP session. This can occur if an application sends non-HTTP sessions using an HTTP destination port.

```
config wanopt profile
  edit default
    config http
      set status enable
      set tunnel-non-http {disable | enable}
    end
```

To drop non-HTTP sessions accepted by the rule set `tunnel-non-http` to `disable`, or set it to `enable` to pass non-HTTP sessions through the tunnel without applying protocol optimization, byte-caching, or web caching. In this case, the FortiGate unit applies TCP protocol optimization to non-HTTP sessions.

Processing unknown HTTP sessions

Unknown HTTP sessions are HTTP sessions that do not comply with HTTP 0.9, 1.0, or 1.1. From the CLI, use the following command to specify how a rule handles such HTTP sessions.

```
config wanopt profile
  edit default
    config http
      set status enable
      set unknown-http-version {best-effort | reject | tunnel}
    end
```

To assume that all HTTP sessions accepted by the rule comply with HTTP 0.9, 1.0, or 1.1, select `best-effort`. If a session uses a different HTTP version, WAN optimization may not parse it correctly. As a result, the

FortiGate unit may stop forwarding the session and the connection may be lost. To reject HTTP sessions that do not use HTTP 0.9, 1.0, or 1.1, select `reject`.

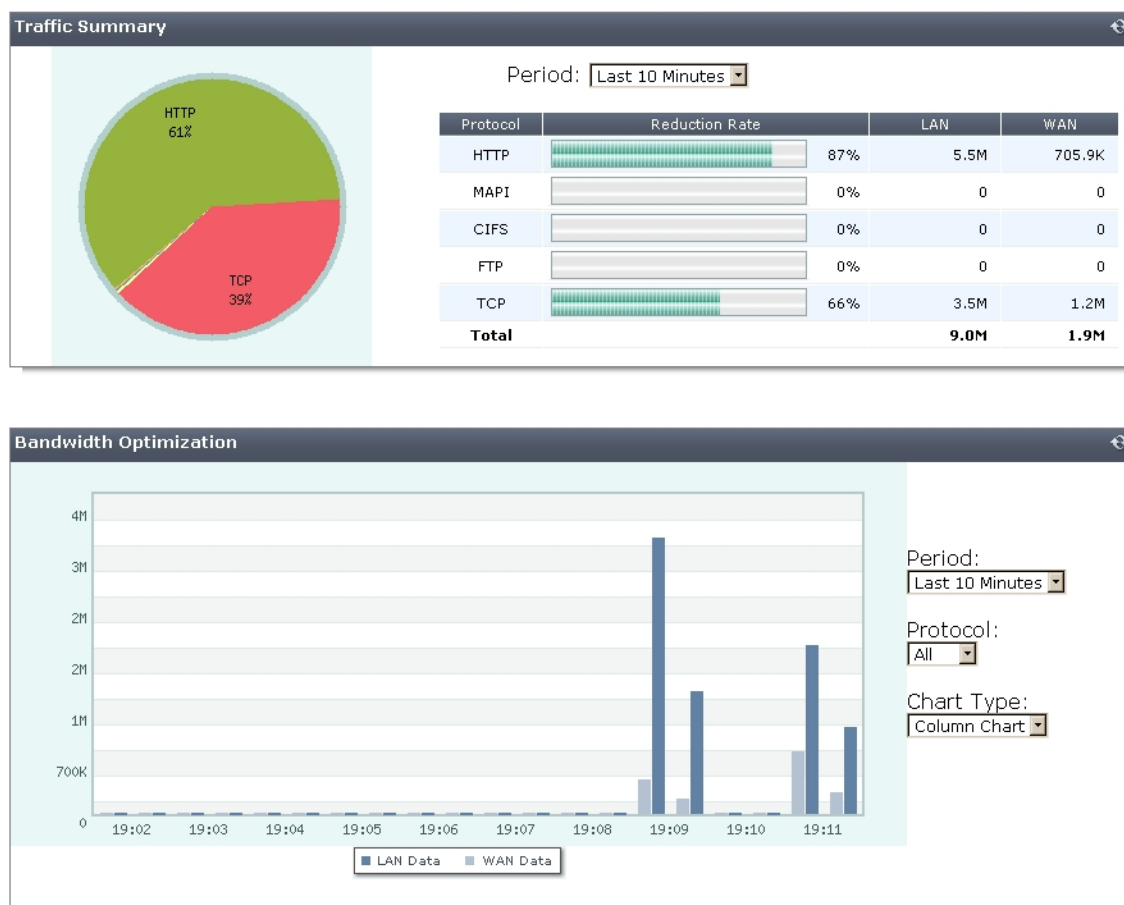
To pass HTTP sessions that do not use HTTP 0.9, 1.0, or 1.1, but without applying HTTP protocol optimization, byte-caching, or web caching, you can also select `tunnel`. TCP protocol optimization is applied to these HTTP sessions.

Monitoring WAN optimization performance

Using WAN optimization monitoring, you can confirm that a FortiGate unit is optimizing traffic and view estimates of the amount of bandwidth saved. The WAN optimization monitor presents collected log information in a graphical format to show network traffic summary and bandwidth optimization information.

To view the WAN optimization monitor, go to **Monitor > WAN Opt. Monitor**.

WAN optimization monitor



Traffic summary

The traffic summary shows how WAN optimization is reducing the amount of traffic on the WAN for each WAN optimization protocol by showing the traffic reduction rate as a percentage of the total traffic. The traffic summary also shows the amount of WAN and LAN traffic. If WAN optimization is being effective the amount of WAN traffic should be lower than the amount of LAN traffic.

You can use the refresh icon to update the traffic summary display at any time. You can also set the amount of time for which the traffic summary shows data. The time period can vary from the last 10 minutes to the last month.

Bandwidth optimization

This section shows network bandwidth optimization per time period. A line or column chart compares an application's pre-optimized size (LAN data) with its optimized size (WAN data). You can select the chart type, the monitoring time period, and the protocol for which to display data. If WAN optimization is being effective the WAN bandwidth should be lower than the LAN bandwidth.

WAN optimization configuration summary

This section includes a client-side and a server-side WAN Optimization configuration summary.:

Client-side configuration summary

WAN optimization profile

Enter the following command to view WAN optimization profile CLI options:

```
tree wanopt profile
-- [profile] --*name (36)
  |- transparent
  |- comments
  |- auth-group (36)
  |- <http> -- status
    |- secure-tunnel
    |- byte-caching
    |- prefer-chunking
    |- tunnel-sharing
    |- log-traffic
    |- port (1,65535)
    |- ssl
    |- ssl-port (1,65535)
    |- unknown-http-version
    +- tunnel-non-http
  |- <cifs> -- status
    |- secure-tunnel
    |- byte-caching
    |- prefer-chunking
    |- tunnel-sharing
    |- log-traffic
    +- port (1,65535)
  |- <mapi> -- status
    |- secure-tunnel
    |- byte-caching
    |- tunnel-sharing
    |- log-traffic
    +- port (1,65535)
  |- <ftp> -- status
    |- secure-tunnel
    |- byte-caching
    |- prefer-chunking
    |- tunnel-sharing
    |- log-traffic
```

```

        +- port (1,65535)
+- <tcp> -- status
    |- secure-tunnel
    |- byte-caching
    |- byte-caching-opt
    |- tunnel-sharing
    |- log-traffic
    |- port
    |- ssl
+- ssl-port (1,65535)

```

Local host ID and peer settings

```

config wanopt settings
    set host-id client
end
config wanopt peer
    edit server
        set ip 10.10.2.82
    end

```

Security policies

Two client-side WAN optimization security policy configurations are possible. One for active-passive WAN optimization and one for manual WAN optimization.

Active/passive mode on the client-side

```

config firewall policy
    edit 2
        set srcintf internal
        set dstintf wan1
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule always
        set service ALL
        set wanopt enable <<< enable WAN optimization
        set wanopt-detection active <<< set the mode to active/passive
        set wanopt-profile "default" <<< select the wanopt profile
    next
end

```

Manual mode on the client-side

```

config firewall policy
    edit 2
        set srcintf internal
        set dstintf wan1
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule always
        set service ALL
        set wanopt enable <<< enable WAN optimization
        set wanopt-detection off <<< sets the mode to manual
        set wanopt-profile "default" <<< select the wanopt profile
    next
end

```

```

        set wanopt-peer "server" <<< set the only peer to do wanopt
                                with
                                (required for manual mode)
    next
end

```

server-side configuration summary

Local host ID and peer settings

```

config wanopt settings
    set host-id server
end
config wanopt peer
    edit client
        set ip 10.10.2.81
    end
end

```

Security policies

Two server-side WAN optimization security policy configurations are possible. One for active-passive WAN optimization and one for manual WAN optimization.

Active/passive mode on server-side

```

config firewall policy
edit 2 <<< the passive mode policy
    set srcintf wan1
    set dstintf internal
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ALL
    set wanopt enable
    set wanopt-detection passive
    set wanopt-passive-opt transparent
end
config firewall proxy-policy
    edit 3 <<< policy that accepts wanopt tunnel connections from the server
        set proxy wanopt <<< wanopt proxy type
        set dstintf internal
        set srcaddr all
        set dstaddr server-subnet
        set action accept
        set schedule always
        set service ALL
    next
end

```

Manual mode on server-side

```

config firewall proxy-policy
    edit 3 <<< policy that accepts wanopt tunnel connections from the client
        set proxy wanopt <<< wanopt proxy type
        set dstintf internal
        set srcaddr all
    end
end

```

```
        set dstaddr server-subnet
        set action accept
        set schedule always
        set service ALL
    next
end
```

WANopt storage

The `config wanopt storage` option has been combined with `config system storage`.

Setting the disk-usage mode is no longer in `config system global`. It is set through `config system storage`.

Syntax:

```
config system storage
    edit <name-string>
        set status enable
        set media-status
        set order
        set partition
        set device
        set size
        set usage
        set wanopt-mode
```

Option	Description
status	Enable/disable storage
media-status	Enable/disable the physical status of current media
order	Set storage order
partition	Label of underlying partition Example: "MIXEDXXE2946380"
device	Partition device. Example: "/dev/vdb1"
size	Partition size. Example: 8616
usage	Use hard disk for logging and WAN Optimization.

Option	Description
	WAN Optimization mode
wanopt-mode	<ul style="list-style-type: none"> • <code>mix</code> - default, recommended • <code>wanopt</code> - recommended if only wanopt feature is enabled • <code>webcache</code> - recommended if only webcache feature is enabled <p>If only one of the two features is being used, using the applicable recommended mode will give a higher cache capacity and improve performance.</p>

WANOpt cache service

The `config wanopt cache-service` command is used to configure `cache-service` clusters between multiple FortiGates. The result is that the `cache-service` daemons of the different FortiGates can collaborate together for serving web cache entries.

To configure the wanopt cache-service

```
config wanopt cache-service
  set prefer-scenario
  set collaboration
  set device-id
  set acceptable-connections
  config dst-peer
    edit <dst-peer-name>
      set auth-type
      set encode-type
      set priority
      set ip
  config src-peer
    edit <src-peer>
      set auth-type
      set encode-type
      set priority
      set ip
```

Option	Description
prefer-scenario	<p>Set the preferred cache behavior to the appropriate balance between latency and hit ratio</p> <p>Options:</p> <ul style="list-style-type: none"> • <code>balance</code> - Balance between speed and cache hit ratio. • <code>prefer-speed</code> - Prefer response speed at the expense of increased cache bypasses. • <code>prefer-cache</code> - Prefer improving hit-ratio through increasing latency tolerance.

Option	Description
collaboration	enable/disable cache collaboration between cache-service clusters
device-id	Set identifier for this cache device
acceptable-connections	Set strategy when accepting cache collaboration connection Options: <ul style="list-style-type: none"> any - The cache-service can accept any cache collaboration connection. peers - The cache-service will only accept connections that are already in src-peers.
auth-type	Set authentication type for this peer Value is integer from 0 to 255
encode-type	Set encode type for this peer Value is integer from 0 to 255
priority	Set priority for this peer Value is integer from 0 to 255. Default = 1
ip	Set cluster IP address of this peer

Video caching

This config wanopt content-delivery-network-rule command configures web-caching including the video-cache matching rules.

To configure the wanopt content-delivery-network-rule

```

config wanopt content-delivery-network-rule
  edit <content_rule_name>
    set comment
    set status
    set host-domain-name-suffix
    set category
    set request-cache-control
    set response-cache-control
    set response-expires
    set text-response-vcache
    set updateserver
    config rules
      edit <rule_name>
        set match-mode
        set skip-rule-mode
        config match-entries
          edit <integer>

```

```

set target
set pattern
config skip-entries
set target
set pattern
config content id
set target
set start-str
set start-skip
set start-direction
set end-str
set end-skip
set end-direction
set range-str

```

Option	Description
comment	Comment about this rule
status	Enable/disable WAN optimization content delivery network rules
host-domain-name-suffix	Suffix portion of the fully qualified domain name (eg. fortinet.com in "www.fortinet.com")
category	Content delivery network rule category
request-cache-control	Enable/disable HTTP request cache control
response-cache-control	Enable/disable HTTP response cache control
response-expires	Enable/disable HTTP response cache expires
updateserver	Enable/disable update server
match-mode	Match criteria for collecting content ID
skip-rule-mode	Skip mode when evaluating skip rules
target	Option in HTTP header or URL parameter to match
pattern	Pattern string for matching target (Referrer or URL pattern, eg. "a", "a*c", "**a*", "a*c*e", and "**")
start-str	String from which to start search
start-skip	Number of characters in URL to skip after <code>start-str</code> has been matched
start-direction	Search direction from <code>start-str</code> match

Option	Description
<code>end-str</code>	String from which to end search
<code>end-skip</code>	Number of characters in URL to skip after end-str has been matched
<code>end-direction</code>	Search direction from <code>end-str</code> match
<code>range-str</code>	Name of content ID within the start string and end string

Best practices

This is a short list of WAN optimization and explicit proxy best practices.

- WAN optimization tunnel sharing is recommended for similar types of WAN optimization traffic. However, tunnel sharing for different types of traffic is not recommended. For example, aggressive and non-aggressive protocols should not share the same tunnel. See [Best practices on page 789](#).
- Active-passive HA is the recommended HA configuration for WAN optimization. See [Best practices on page 789](#).
- Configure WAN optimization authentication with specific peers. Accepting any peer is not recommended as this can be less secure. See [Accepting any peers on page 1](#).
- Set the explicit proxy **Default Firewall Policy Action** to **Deny**. This means that a security policy is required to use the explicit web proxy. See [General explicit web proxy configuration steps on page 1](#).
- Set the explicit FTP proxy **Default Firewall Policy Action** to **Deny**. This means that a security policy is required to use the explicit FTP proxy. See [General explicit FTP proxy configuration steps on page 1](#).
- Do not enable the explicit web or FTP proxy on an interface connected to the Internet. This is a security risk because anyone on the Internet who finds the proxy could use it to hide their source address. If you must enable the proxy on such an interface make sure authentication is required to use the proxy. See [General explicit web proxy configuration steps on page 1](#).

Example basic manual (peer-to-peer) WAN optimization configuration

In a manual (peer to peer) configuration the WAN optimization tunnel can be set up between one client-side FortiGate unit and one server-side FortiGate unit. The peer ID of the server-side FortiGate unit is added to the client-side WAN optimization policy. When the client-side FortiGate unit initiates a tunnel with the server-side FortiGate unit, the packets that initiate the tunnel include information that allows the server-side FortiGate unit to determine that it is a manual tunnel request. The server-side FortiGate unit does not require a WAN optimization profile; you just need to add the client peer host ID and IP address to the server-side FortiGate unit peer list and from the CLI an explicit proxy policy to accept WAN optimization tunnel connections.

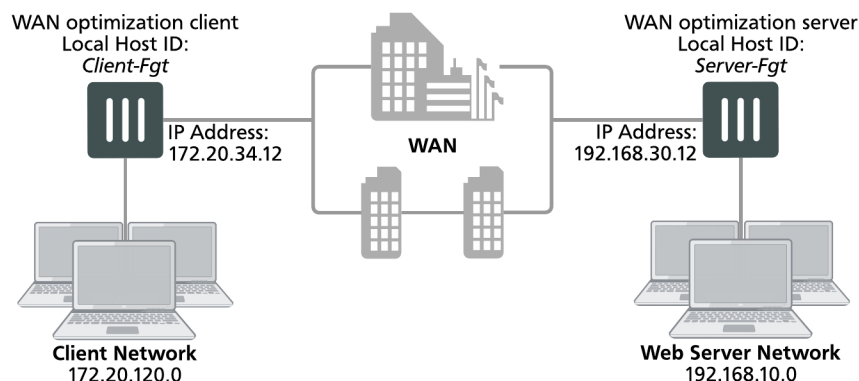
In a manual WAN optimization configuration, you create a manual WAN optimization security policy on the client-side FortiGate unit. To do this you must use the CLI to set `wanopt-detection` to `off` and to add the peer host ID of the server-side FortiGate unit to the WAN optimization security policy.

Network topology and assumptions

This example configuration includes a client-side FortiGate unit called Client-Fgt with a WAN IP address of 172.20.34.12. This unit is in front of a network with IP address 172.20.120.0. The server-side FortiGate unit is called Server_Fgt with a WAN IP address of 192.168.30.12. This unit is in front of a web server network with IP address 192.168.10.0.

This example customizes the default WAN optimization profile on the client-side FortiGate unit and adds it to the WAN optimization policy. You can also create a new WAN optimization profile.

Example manual (peer-to-peer) topology



General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Configure the client-side FortiGate unit:
 - Add peers.
 - Configure the default WAN optimization profile to optimize HTTP traffic.
 - Add a manual WAN optimization security policy.
2. Configure the server-side FortiGate unit:
 - Add peers.
 - Add a WAN optimization tunnel policy.

Configuring basic peer-to-peer WAN optimization - web-based manager

Use the following steps to configure the example configuration from the web-based manager.

To configure the client-side FortiGate unit

1. Go to **WAN Opt. & Cache > Peers** and enter a **Local Host ID** for the client-side FortiGate unit:

Local Host ID	Client-Fgt
----------------------	------------

2. Select **Apply**.
3. Select **Create New** and add the server-side FortiGate unit **Peer Host ID** and **IP Address** for the server-side FortiGate:

Peer Host ID	Server-Fgt
IP Address	192.168.30.12

4. Select **OK**.
5. Go to **Policy & Objects > Addresses** and select **Create New** to add a firewall address for the client network.

Category	Address
Name	Client-Net
Type	Subnet
Subnet / IP Range	172.20.120.0/24
Interface	port1

6. Select **Create New** to add a firewall address for the web server network.

Category	Address
Name	Web-Server-Net
Type	Subnet
Subnet / IP Range	192.168.10.0/24
Interface	port2

7. Go to **WAN Opt. & Cache > Profiles** and edit the default profile.
8. Select **Transparent Mode**.
9. Under Protocol, select **HTTP** and for HTTP select **Byte Caching**. Leave the HTTP **Port** set to 80.
10. Select **Apply** to save your changes.
11. Go to **Policy & Objects > IPv4 Policy** and add a WAN optimization security policy to the client-side FortiGate unit that accepts traffic to be optimized:

Incoming Interface	port1
Source Address	all
Outgoing Interface	port2
Destination Address	all
Schedule	always
Service	ALL
Action	ACCEPT

12. Select **Enable WAN Optimization** and configure the following settings:

Enable WAN Optimization	active
Profile	default

13. Select **OK**.
14. Edit the policy from the CLI to turn off `wanopt-detection`, add the peer ID of the server-side FortiGate unit, and the default WAN optimization profile. The following example assumes the ID of the policy is 5:

```
config firewall policy
edit 5
    set wanopt-detection off
    set wanopt-peer Server-Fgt
    set wanopt-profile default
end
```

When you set the detection mode to `off` the policy becomes a manual mode WAN optimization policy. On the web-based manager the WAN optimization part of the policy changes to the following:

Enable WAN Optimization	Manual (Profile: default, Peer: Peer-Fgt-2)
--------------------------------	---

To configure the server-side FortiGate unit

1. Go to **WAN Opt. & Cache > Peers** and enter a **Local Host ID** for the server-side FortiGate unit:

Local Host ID	Server-Fgt
----------------------	------------

2. Select **Apply**.
3. Select **Create New** and add a **Peer Host ID** and the **IP Address** for the client-side FortiGate unit:

Peer Host ID	Client-Fgt
IP Address	172.20.34.12

4. Select **OK**.
5. Enter the following CLI command to add an explicit proxy policy to accept WAN optimization tunnel connections.

```
configure firewall proxy-policy
edit 0
    set proxy wanopt
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ALL
next
end
```

Configuring basic peer-to-peer WAN optimization - CLI

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit CLI.

To configure the client-side FortiGate unit

1. Add the Local Host ID to the client-side FortiGate configuration:

```
config wanopt settings
    set host-id Client-Fgt
end
```

2. Add the server-side Local Host ID to the client-side peer list:

```
config wanopt peer
  edit Server-Fgt
    set ip 192.168.30.12
  end
```

3. Add a firewall address for the client network.

```
config firewall address
  edit Client-Net
    set type ipmask
    set subnet 172.20.120.0 255.255.255.0
    set associated-interface port1
  end
```

4. Add a firewall address for the web server network.

```
config firewall address
  edit Web-Server-Net
    set type ipmask
    set subnet 192.168.10.0 255.255.255.0
    set associated-interface port2
  end
```

5. Edit the default WAN optimization profile, select transparent mode, enable HTTP WAN optimization and enable byte caching for HTTP. Leave the HTTP Port set to 80.

```
config wanopt profile
  edit default
    set transparent enable
    config http
      set status enable
      set byte-caching enable
    end
  end
```

6. Add a WAN optimization security policy to the client-side FortiGate unit to accept the traffic to be optimized:

```
config firewall policy
  edit 0
    set srcintf port1
    set dstintf port2
    set srcaddr all
    set dstaddr all
    set action accept
    set service ALL
    set schedule always
    set wanopt enable
    set wanopt-profile default
    set wanopt-detection off
    set wanopt-peer Server-Fgt
  end
```

To configure the server-side FortiGate unit**1. Add the Local Host ID to the server-side FortiGate configuration:**

```
config wanopt settings
  set host-id Server-Fgt
```



```
end
```

2. Add the client-side Local Host ID to the server-side peer list:

```
config wanopt peer
  edit Client-Fgt
    set ip 192.168.30.12
  end
```

3. Add a WAN optimization tunnel explicit proxy policy.

```
configure firewall proxy-policy
  edit 0
    set proxy wanopt
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ALL
  next
end
```

Testing and troubleshooting the configuration

To test the configuration attempt to start a web browsing session between the client network and the web server network. For example, from a PC on the client network browse to the IP address of a web server on the web server network, for example `http://192.168.10.100`. Even though this address is not on the client network you should be able to connect to this web server over the WAN optimization tunnel.

If you can connect, check WAN optimization monitoring. If WAN optimization has been forwarding the traffic the WAN optimization monitor should show the protocol that has been optimized (in this case HTTP) and the reduction rate in WAN bandwidth usage.

If you can't connect you can try the following to diagnose the problem:

- Review your configuration and make sure all details such as address ranges, peer names, and IP addresses are correct.
- Confirm that the security policy on the client-side FortiGate unit is accepting traffic for the 192.168.10.0 network. You can do this by checking the policy monitor (**Monitor > Firewall User Monitor**). Look for sessions that use the policy ID of this policy.
- Check routing on the FortiGate units and on the client and web server networks to make sure packets can be forwarded as required. The FortiGate units must be able to communicate with each other, routing on the client network must allow packets destined for the web server network to be received by the client-side FortiGate unit, and packets from the server-side FortiGate unit must be able to reach the web servers.

You can use the following `get` and `diagnose` commands to display information about how WAN optimization is operating.

Enter the following command to list all of the running WAN optimization tunnels and display information about each one. The command output for the client-side FortiGate unit shows 10 tunnels all created by peer-to-peer WAN optimization rules (auto-detect set to off).

```
diagnose wad tunnel list

Tunnel: id=100 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web-servers id=100 ip=192.168.30.12
```

```
SSL-secured-tunnel=no auth-grp=
bytes_in=348 bytes_out=384

Tunnel: id=99 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web-servers id=99 ip=192.168.30.12
SSL-secured-tunnel=no auth-grp=
bytes_in=348 bytes_out=384

Tunnel: id=98 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web-servers id=98 ip=192.168.30.12
SSL-secured-tunnel=no auth-grp=
bytes_in=348 bytes_out=384

Tunnel: id=39 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web-servers id=39 ip=192.168.30.12
SSL-secured-tunnel=no auth-grp=
bytes_in=1068 bytes_out=1104

Tunnel: id=7 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web-servers id=7 ip=192.168.30.12
SSL-secured-tunnel=no auth-grp=
bytes_in=1228 bytes_out=1264

Tunnel: id=8 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web-servers id=8 ip=192.168.30.12
SSL-secured-tunnel=no auth-grp=
bytes_in=1228 bytes_out=1264

Tunnel: id=5 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web-servers id=5 ip=192.168.30.12
SSL-secured-tunnel=no auth-grp=
bytes_in=1228 bytes_out=1264

Tunnel: id=4 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web-servers id=4 ip=192.168.30.12
SSL-secured-tunnel=no auth-grp=
bytes_in=1228 bytes_out=1264

Tunnel: id=1 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web-servers id=1 ip=192.168.30.12
SSL-secured-tunnel=no auth-grp=
bytes_in=1228 bytes_out=1264

Tunnel: id=2 type=manual
vd=0 shared=no uses=0 state=3
peer name=Web-servers id=2 ip=192.168.30.12
SSL-secured-tunnel=no auth-grp=
bytes_in=1228 bytes_out=1264
```

```
Tunnels total=10 manual=10 auto=0
```

Example active-passive WAN optimization

In active-passive WAN optimization you add an active WAN optimization policy to the client-side FortiGate unit and you add a WAN optimization tunnel policy and a passive WAN optimization policy to the server-side FortiGate unit.

The active policy accepts the traffic to be optimized and sends it down the WAN optimization tunnel to the server-side FortiGate unit. The active policy can also apply security profiles and other features to traffic before it exits the client-side FortiGate unit.

A tunnel explicit proxy policy on the server-side FortiGate unit allows the server-side FortiGate unit to form a WAN optimization tunnel with the client-side FortiGate unit. The passive WAN optimization policy is required because of the active policy on the client-side FortiGate unit. You can also use the passive policy to apply WAN optimization transparent mode and features such as security profiles, logging, traffic shaping and web caching to the traffic before it exits the server-side FortiGate unit.

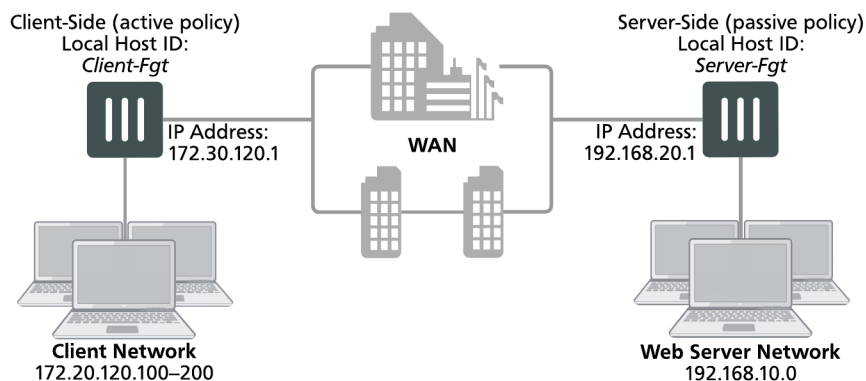
Network topology and assumptions

On the client-side FortiGate unit this example configuration includes a WAN optimization profile that optimizes CIFS, HTTP, and FTP traffic and an active WAN optimization policy. The active policy also applies virus scanning to the WAN optimization traffic.

On the server-side FortiGate unit, the passive policy applies application control to the WAN optimization traffic.

In this example, WAN optimization transparent mode is selected in the WAN optimization profile and the passive WAN optimization policy accepts this transparent mode setting. This means that the optimized packets maintain their original source and destination addresses. As a result, routing on the client network must be configured to route packets for the server network to the client-side FortiGate unit. Also the routing configuration on the server network must be able to route packets for the client network to the server-side FortiGate unit.

Example active-passive WAN optimization topology



General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Configure the client-side FortiGate unit:

- Add peers.
 - Add a WAN optimization profile to optimize CIFS, FTP, and HTTP traffic.
 - Add firewall addresses for the client and web server networks.
 - Add an active WAN optimization policy.
2. Configure the server-side FortiGate unit by:
 - Add peers.
 - Add firewall addresses for the client and web server networks.
 - Add a passive WAN optimization policy.
 - Add a WAN optimization tunnel policy.

Configuring basic active-passive WAN optimization - web-based manager

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit web-based manager.

To configure the client-side FortiGate unit

1. Go to **WAN Opt. & Cache > Peers** and enter a **Local Host ID** for the client-side FortiGate unit:

Local Host ID	Client-Fgt
----------------------	------------

2. Select **Apply**.
3. Select **Create New** and add a Peer Host ID and the **IP Address** for the server-side FortiGate unit:

Peer Host ID	Server-Fgt
IP Address	192.168.20.1

4. Select **OK**.
5. Go to **WAN Opt. & Cache > Profiles** and select **Create New** to add a WAN optimization profile to optimize CIFS, HTTP, and FTP traffic:

Name	Custom-wan-opt-pro
Transparent Mode	Select

6. Select the **CIFS** protocol, select **Byte Caching** and set the **Port** to 445.
7. Select the **FTP** protocol, select **Byte Caching** and set the **Port** to 21.
8. Select the **HTTP** protocol, select **Byte Caching** and set the **Port** to 80.
9. Select **OK**.
10. Go to **Policy & Objects > Addresses** and select **Create New** to add an address for the client network.

Category	Address
Address Name	Client-Net
Type	IP Range

Subnet / IP Range	172.20.120.100-172.20.120.200
Interface	port1

11. Select **Create New** to add an address for the web server network.

Category	Address
Address Name	Web-Server-Net
Type	Subnet
Subnet / IP Range	192.168.10.0/24
Interface	port2

12. Go to **Policy & Objects > IPv4 Policy** and select **Create New** to add an active WAN optimization security policy:

Incoming Interface	port1
Source Address	Client-Net
Outgoing Interface	port2
Destination Address	Web-Server-Net
Schedule	always
Service	HTTP FTP SMB
Action	ACCEPT

13. Turn on **WAN Optimization** and configure the following settings:

WAN Optimization	active
Profile	Custom-wan-opt-pro

14. Turn on Antivirus and select the **default** antivirus profile.
15. Select **OK**.

To configure the server-side FortiGate unit

1. Go to **WAN Opt. & Cache > Peers** and enter a **Local Host ID** for the server-side FortiGate unit:

Local Host ID	Server-Fgt
----------------------	------------

2. Select **Apply**.
3. Select **Create New** and add a **Peer Host ID** and the **IP Address** for the client-side FortiGate unit:

Peer Host ID	Client-Fgt
IP Address	172.30.120.1

4. Select **OK**.
5. Go to **Policy & Objects > Addresses** and select **Create New** to add an address for the client network.

Category	Address
Address Name	Client-Net
Type	IP Range
Subnet / IP Range	172.20.120.100-172.20.120.200
Interface	port1

6. Select **Create New** to add a firewall address for the web server network.

Category	Address
Address Name	Web-Server-Net
Type	Subnet
Subnet / IP Range	192.168.10.0/24
Interface	port2

7. Select **OK**.
8. Select **Policy & Objects > IPv4 Policy** and select **Create New** to add a passive WAN optimization policy that applies application control.

Incoming Interface	port2
Source Address	Client-Net
Outgoing Interface	port1
Destination Address	Web-Server-Net
Schedule	always
Service	ALL
Action	ACCEPT

9. Turn on **WAN Optimization** and configure the following settings:

WAN Optimization	passive
Passive Option	default

10. Select **OK**.

11. From the CLI enter the following command to add a WAN optimization tunnel explicit proxy policy.

```
configure firewall proxy-policy
edit 0
    set proxy wanopt
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ALL
next
end
```

Configuring basic active-passive WAN optimization - CLI

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit CLI.

To configure the client-side FortiGate unit

1. Add the Local Host ID to the client-side FortiGate configuration:

```
config wanopt settings
    set host-id Client-Fgt
end
```

2. Add the server-side Local Host ID to the client-side peer list:

```
config wanopt peer
    edit Server-Fgt
        set ip 192.168.20.1
    end
end
```

3. Add a WAN optimization profile to optimize CIFS, HTTP, and FTP traffic.

```
config wanopt profile
    edit Custom-wan-opt-pro
        config cifs
            set status enable
            set byte-caching enable
            set port 445
        end
        config http
            set status enable
            set byte-caching enable
            set port 80
        end
        config ftp
            set status enable
            set byte-caching enable
            set port 21
        end
    end
end
```

4. Add a firewall address for the client network.

```
config firewall address
    edit Client-Net
        set type iprange
        set start-ip 172.20.120.100
        set end-ip 172.20.120.200
    end
end
```

```

        set associated-interface port1
    end

```

5. Add a firewall address for the web server network.

```

config firewall address
    edit Web-Server-Net
        set type ipmask
        set subnet 192.168.10.0 255.255.255.0
        set associated-interface port2
    end

```

6. Add an active WAN optimization security policy that applies virus scanning:

```

config firewall policy
    edit 0
        set srcintf port1
        set dstintf port2
        set srcaddr Client-net
        set dstaddr Web-Server-Net
        set action accept
        set service HTTP FTP SMB
        set schedule always
        set wanopt enable
        set wanopt-detection active
        set wanopt-profile Custom-wan-opt-pro
    end

```

To configure the server-side FortiGate unit

1. Add the Local Host ID to the server-side FortiGate configuration:

```

config wanopt settings
    set host-id Server-Fgt
end

```

2. Add the client-side Local Host ID to the server-side peer list:

```

config wanopt peer
    edit Client-Fgt
        set ip 172.20.120.1
    end

```

3. Add a firewall address for the client network.

```

config firewall address
    edit Client-Net
        set type iprange
        set start-ip 172.20.120.100
        set end-ip 172.20.120.200
        set associated-interface port1
    end

```

4. Add a firewall address for the web server network.

```

config firewall address
    edit Web-Server-Net
        set type ipmask
        set subnet 192.168.10.0 255.255.255.0
        set associated-interface port2
    end

```

5. Add a passive WAN optimization policy.

```

config firewall policy
    edit 0
        set srcintf port1
        set dstintf port2

```



```

        set srcaddr Client-Net
        set dstaddr Web-Server-Net
        set action accept
        set service ALL
        set schedule always
        set wanopt enable
        set wanopt-detection passive
        set wanopt-passive-opt default
    end

```

6. Add a WAN optimization tunnel explicit proxy policy.

```

configure firewall proxy-policy
edit 0
    set proxy wanopt
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ALL
next
end

```

Testing and troubleshooting the configuration

To test the configuration attempt to start a web browsing session between the client network and the web server network. For example, from a PC on the client network browse to the IP address of a web server on the web server network, for example `http://192.168.10.100`. Even though this address is not on the client network you should be able to connect to this web server over the WAN optimization tunnel.

If you can connect, check WAN optimization monitoring. If WAN optimization has been forwarding the traffic the WAN optimization monitor should show the protocol that has been optimized (in this case HTTP) and the reduction rate in WAN bandwidth usage.

If you can't connect you can try the following to diagnose the problem:

- Review your configuration and make sure all details such as address ranges, peer names, and IP addresses are correct.
- Confirm that the security policy on the Client-Side FortiGate unit is accepting traffic for the 192.168.10.0 network and that this security policy does not include security profiles. You can do this by checking the FortiGate session table from the dashboard. Look for sessions that use the policy ID of this policy.
- Check routing on the FortiGate units and on the client and web server networks to make sure packets can be forwarded as required. The FortiGate units must be able to communicate with each other, routing on the client network must allow packets destined for the web server network to be received by the client-side FortiGate unit, and packets from the server-side FortiGate unit must be able to reach the web servers etc.

You can use the following `get` and `diagnose` commands to display information about how WAN optimization is operating

Enter the following command to list all of the running WAN optimization tunnels and display information about each one. The command output shows 3 tunnels all created by peer-to-peer WAN optimization rules (auto-detect set to on).

```

diagnose wad tunnel list

Tunnel: id=139 type=auto

```

```

vd=0 shared=no uses=0 state=1
peer name= id=0 ip=unknown
SSL-secured-tunnel=no auth-grp=test
bytes_in=744 bytes_out=76

Tunnel: id=141 type=auto
vd=0 shared=no uses=0 state=1
peer name= id=0 ip=unknown
SSL-secured-tunnel=no auth-grp=test
bytes_in=727 bytes_out=76

Tunnel: id=142 type=auto
vd=0 shared=no uses=0 state=1
peer name= id=0 ip=unknown
SSL-secured-tunnel=no auth-grp=test
bytes_in=727 bytes_out=76

Tunnels total=3 manual=0 auto=3

```

Example adding secure tunneling to an active-passive WAN optimization configuration

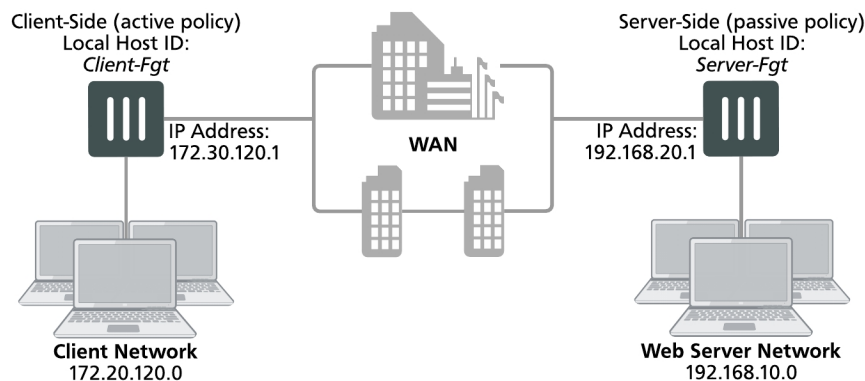
This example shows how to configure two FortiGate units for active-passive WAN optimization with secure tunneling. The same authentication group is added to both FortiGate units. The authentication group includes a password (or pre-shared key) and has **Peer Acceptance** set to **Accept any Peer**. An active policy is added to the client-side FortiGate unit and a passive policy to the server-side FortiGate unit. The active policy includes a profile that performs secure tunneling, optimizes HTTP traffic, and uses transparent mode and byte caching.

The authentication group is named **Auth-Secure-Tunnel** and the password for the pre-shared key is **2345678**. The topology for this example is shown below. This example includes web-based manager configuration steps followed by equivalent CLI configuration steps. For information about secure tunneling, see [Secure tunneling on page 1](#).

Network topology and assumptions

This example configuration includes a client-side FortiGate unit called Client-net with a WAN IP address of 172.30.120.1. This unit is in front of a network with IP address 172.20.120.0. The server-side FortiGate unit is called Web-servers and has a WAN IP address of 192.168.20.1. This unit is in front of a web server network with IP address 192.168.10.0.

Example active-passive WAN optimization and secure tunneling topology



General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Configure the client-side FortiGate unit:
 - Add peers.
 - Add an authentication group.
 - Add an active WAN optimization policy.
2. Configure the server-side FortiGate unit.
 - Add peers.
 - Add the same authentication group
 - Add a passive WAN optimization policy that applies application control.
 - Add a WAN optimization tunnel policy.

Also note that if you perform any additional actions between procedures, your configuration may have different results.

Configuring WAN optimization with secure tunneling - web-based manager

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit web-based manager. (CLI steps follow.)

To configure the client-side FortiGate unit

1. Go to **WAN Opt. & Cache > Peers** and enter a **Local Host ID** for the client-side FortiGate unit:

Local Host ID	Client-Fgt
----------------------	------------

2. Select **Apply** to save your setting.
3. Select **Create New** and add a **Peer Host ID** and the **IP Address** for the server-side FortiGate unit:

Peer Host ID	Server-Fgt
IP Address	192.168.20.1

4. Select **OK**.
5. Go to **WAN Opt. & Cache > Authentication Groups** and select **Create New** to add the authentication group to be used for secure tunneling:

Name	Auth-Secure-Tunnel
Authentication Method	Pre-shared key
Password	2345678
Peer Acceptance	Accept Any Peer

6. Select **OK**.
7. Go to **WAN Opt. & Cache > Profiles** and select **Create New** to add a WAN optimization profile that enables secure tunneling and includes the authentication group:

Name	Secure-wan-op-pro
Transparent Mode	Select
Authentication Group	Auth-Secure-tunnel

8. Select the **HTTP** protocol, select Secure Tunneling and **Byte Caching** and set the **Port** to 80.
9. Select **OK**.
10. Go to **Policy & Objects > Addresses** and select **Create New** to add a firewall address for the client network.

Category	Address
Name	Client-Net
Type	Subnet
Subnet / IP Range	172.20.120.0/24
Interface	port1

11. Select **Create New** to add a firewall address for the web server network.

Category	Address
Address Name	Web-Server-Net
Type	Subnet
Subnet / IP Range	192.168.10.0/24
Interface	port2

12. Go to **Policy & Objects > IPv4 Policy** and select **Create New** to add an active WAN optimization security policy:

Incoming Interface	port1
Source Address	Client-Net
Outgoing Interface	port2
Destination Address	Web-Server-Net
Schedule	always
Service	HTTP
Action	ACCEPT

13. Turn on **WAN Optimization** and configure the following settings:

WAN Optimization	active
Profile	Secure-wan-opt-pro

14. Select **OK**.

To configure the server-side FortiGate unit

1. Go to **WAN Opt. & Cache > Peers** and enter a **Local Host ID** for the server-side FortiGate unit:

Local Host ID	Server-Fgt
----------------------	------------

2. Select **Apply** to save your setting.
 3. Select **Create New** and add a **Peer Host ID** and the **IP Address** for the client-side FortiGate unit:

Peer Host ID	Client-Fgt
IP Address	172.30.120.1

4. Select **OK**.
 5. Go to **WAN Opt. & Cache > Authentication Groups** and select **Create New** and add an authentication group to be used for secure tunneling:

Name	Auth-Secure-Tunnel
Authentication Method	Pre-shared key
Password	2345678
Peer Acceptance	Accept Any Peer

6. Select **OK**.
 7. Go to **Policy & Objects > Addresses** and select **Create New** to add a firewall address for the client network.

Category	Address
Name	Client-Net
Type	Subnet
Subnet / IP Range	172.20.120.0/24
Interface	port1

8. Select **Create New** to add a firewall address for the web server network.

Category	Address
-----------------	---------

Address Name	Web-Server-Net
Type	Subnet
Subnet / IP Range	192.168.10.0/24
Interface	port2

9. Select **OK**.
10. Select **Create New** to add a passive WAN optimization policy that applies application control.

Incoming Interface	port2
Source Address	Client-Net
Outgoing Interface	port1
Destination Address	Web-Server-Net
Schedule	always
Service	ALL
Action	ACCEPT

11. Turn on **WAN Optimization** and configure the following settings:

WAN Optimization	passive
Passive Option	default

12. Select **OK**.
13. From the CLI enter the following command to add a WAN optimization tunnel explicit proxy policy.

```
configure firewall proxy-policy
edit 0
set proxy wanopt
set dstintf port1
set srcaddr all
set dstaddr all
set action accept
set schedule always
set service ALL
next
end
```

Configuring WAN optimization with secure tunneling - CLI

Use the following steps to configure the example WAN optimization configuration from the client-side and server-side FortiGate unit CLI.

To the client-side FortiGate unit

1. Add the Local Host ID to the client-side FortiGate configuration:

```
config wanopt settings
```

```
set host-id Client-Fgt
end
```

2. Add the server-side Local Host ID to the client-side peer list:

```
config wanopt peer
edit Server-Fgt
set ip 192.168.20.1
end
```

3. Add a new authentication group to be used for secure tunneling:

```
config wanopt auth-group
edit Auth-Secure-Tunnel
set auth-method psk
set psk 2345678
end
```

Leave peer-accept at its default value.

4. Add a WAN optimization profile that enables secure tunneling and includes the authentication group, enables HTTP protocol optimization, and enables secure tunneling and byte caching for HTTP traffic:

```
config wanopt profile
edit Secure-wan-op-pro
set auth-group Auth-Secure-Tunnel
config http
set status enable
set secure-tunnel enable
set byte-caching enable
set port 80
end
end
```

5. Add a firewall address for the client network.

```
config firewall address
edit Client-Net
set type ipmask
set subnet 172.20.120.0 255.255.255.0
set associated-interface port1
end
```

6. Add a firewall address for the web server network.

```
config firewall address
edit Web-Server-Net
set type ipmask
set subnet 192.168.10.0 255.255.255.0
set associated-interface port2
end
```

7. Add an active WAN optimization security policy that includes the WAN optimization profile that enables secure tunneling and that applies virus scanning:

```
config firewall policy
edit 0
set srcintf port1
set dstintf port2
set srcaddr Client-Net
set dstaddr Web-Server-Net
set action accept
set service HTTP
```

```
set schedule always
set wanopt enable
set wanopt-detection active
set wanopt-profile Secure-wan-opt-pro
end
```

To configure the server-side FortiGate unit

1. Add the Local Host ID to the server-side FortiGate configuration:

```
config wanopt settings
set host-id Server-Fgt
end
```

2. Add the client-side Local Host ID to the server-side peer list:

```
config wanopt peer
edit Client-Fgt
set ip 172.20.120.1
end
```

3. Add an authentication group to be used for secure tunneling:

```
config wanopt auth-group
edit Auth-Secure-Tunnel
set auth-method psk
set psk 2345678
end
```

Leave `peer-accept` at its default value.

4. Add a firewall address for the client network.

```
config firewall address
edit Client-Net
set type ipmask
set subnet 172.20.120.0 255.255.255.0
set associated-interface port1
end
```

5. Add a firewall address for the web server network.

```
config firewall address
edit Web-Server-Net
set type ipmask
set subnet 192.168.10.0 255.255.255.0
set associated-interface port2
end
```

6. Add a passive WAN optimization policy.

```
config firewall policy
edit 0
set srcintf port1
set dstintf port2
set srcaddr Client-Net
set dstaddr Web-Server-Net
set action accept
set service ALL
set schedule always
set wanopt enable
set wanopt-detection passive
```



```
    set wanopt-passive-opt default
end
```

7. Add a WAN optimization tunnel explicit proxy policy.

```
configure firewall proxy-policy
edit 0
    set proxy wanopt
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ALL
next
end
```

Peers and authentication groups

All communication between WAN optimization peers begins with one WAN optimization peer (or client-side FortiGate unit) sending a WAN optimization tunnel request to another peer (or server-side FortiGate unit). During this process, the WAN optimization peers identify and optionally authenticate each other.

Basic WAN optimization peer requirements

WAN optimization requires the following configuration on each peer. For information about configuring local and peer host IDs, see [Basic WAN optimization peer requirements on page 811](#).

- The peer must have a unique host ID.
- Unless authentication groups are used, peers authenticate each other using host ID values. Do not leave the local host ID at its default value.
- The peer must know the host IDs and IP addresses of all of the other peers that it can start WAN optimization tunnels with. This does not apply if you use authentication groups that accept all peers.
- All peers must have the same local certificate installed on their FortiGate units if the units authenticate by local certificate. Similarly, if the units authenticate by pre-shared key (password), administrators must know the password. The type of authentication is selected in the authentication group. This applies only if you use authentication groups.

Accepting any peers

Strictly speaking, you do not need to add peers. Instead you can configure authentication groups that accept any peer. However, for this to work, both peers must have the same authentication group (with the same name) and both peers must have the same certificate or pre-shared key.

Accepting any peer is useful if you have many peers or if peer IP addresses change. For example, you could have FortiGate units with dynamic external IP addresses (using DHCP or PPPoE). For most other situations, this method is not recommended and is not a best practice as it is less secure than accepting defined peers or a single peer. For more information, see [Basic WAN optimization peer requirements on page 811](#).

How FortiGate units process tunnel requests for peer authentication

When a client-side FortiGate unit attempts to start a WAN optimization tunnel with a peer server-side FortiGate unit, the tunnel request includes the following information:

- the client-side local host ID
- the name of an authentication group, if included in the rule that initiates the tunnel
- if an authentication group is used, the authentication method it specifies: pre-shared key or certificate
- the type of tunnel (secure or not).

For information about configuring the local host ID, peers and authentication groups, see [How FortiGate units process tunnel requests for peer authentication on page 811](#) and [How FortiGate units process tunnel requests for peer authentication on page 811](#).

The authentication group is optional unless the tunnel is a secure tunnel. For more information, see [How FortiGate units process tunnel requests for peer authentication on page 811](#).

If the tunnel request includes an authentication group, the authentication will be based on the settings of this group as follows:

- The server-side FortiGate unit searches its own configuration for the name of the authentication group in the tunnel request. If no match is found, the authentication fails.
- If a match is found, the server-side FortiGate unit compares the authentication method in the client and server authentication groups. If the methods do not match, the authentication fails.
- If the authentication methods match, the server-side FortiGate unit tests the peer acceptance settings in its copy of the authentication group.
- If the setting is **Accept Any Peer**, the authentication is successful.
- If the setting is **Specify Peer**, the server-side FortiGate unit compares the client-side local host ID in the tunnel request with the peer name in the server-side authentication group. If the names match, authentication is successful. If a match is not found, authentication fails.
- If the setting is **Accept Defined Peers**, the server-side FortiGate unit compares the client-side local host ID in the tunnel request with the server-side peer list. If a match is found, authentication is successful. If a match is not found, authentication fails.

If the tunnel request does not include an authentication group, authentication will be based on the client-side local host ID in the tunnel request. The server-side FortiGate unit searches its peer list to match the client-side local host ID in the tunnel request. If a match is found, authentication is successful. If a match is not found, authentication fails.

If the server-side FortiGate unit successfully authenticates the tunnel request, the server-side FortiGate unit sends back a tunnel setup response message. This message includes the server-side local host ID and the authentication group that matches the one in the tunnel request.

The client-side FortiGate unit then performs the same authentication procedure as the server-side FortiGate unit did. If both sides succeed, tunnel setup continues.

Configuring peers

When you configure peers, you first need to add the local host ID that identifies the FortiGate unit for WAN optimization and then add the peer host ID and IP address of each FortiGate unit with which a FortiGate unit can create WAN optimization tunnels.

To configure WAN optimization peers - web-based manager:

1. Go to **WAN Opt. & Cache > Peers**.
2. For **Local Host ID**, enter the local host ID of **this** FortiGate unit and select **Apply**. If you add this FortiGate unit as a peer to another FortiGate unit, use this ID as its **peer** host ID.

The local or host ID can contain up to 25 characters and can include spaces.

3. Select **Create New** to add a new peer.
4. For **Peer Host ID**, enter the peer host ID of the peer FortiGate unit. This is the local host ID added to the peer FortiGate unit.
5. For **IP Address**, add the IP address of the peer FortiGate unit. This is the source IP address of tunnel requests sent by the peer, usually the IP address of the FortiGate interface connected to the WAN.
6. Select **OK**.

To configure WAN optimization peers - CLI:

In this example, the local host ID is named `HQ_Peer` and has an IP address of `172.20.120.100`. Three peers are added, but you can add any number of peers that are on the WAN.

1. Enter the following command to set the local host ID to `HQ_Peer`.

```
config wanopt settings
  set host-id HQ_peer
end
```

2. Enter the following commands to add three peers.

```
config wanopt peer
  edit Wan_opt_peer_1
    set ip 172.20.120.100
  next
  edit Wan_opt_peer_2
    set ip 172.30.120.100
  next
  edit Wan_opt_peer_3
    set ip 172.40.120.100
end
```

Configuring authentication groups

You need to add authentication groups to support authentication and secure tunneling between WAN optimization peers.

To perform authentication, WAN optimization peers use a certificate or a pre-shared key added to an authentication group so they can identify each other before forming a WAN optimization tunnel. Both peers must have an authentication group with the same name and settings. You add the authentication group to a peer-to-peer or active rule on the client-side FortiGate unit. When the server-side FortiGate unit receives a tunnel start request from the client-side FortiGate unit that includes an authentication group, the server-side FortiGate unit finds an authentication group in its configuration with the same name. If both authentication groups have the same certificate or pre-shared key, the peers can authenticate and set up the tunnel.

Authentication groups are also required for secure tunneling.

To add authentication groups, go to **WAN Opt. & Cache > Authentication Groups**.

To add an authentication group - web-based manager:

Use the following steps to add any kind of authentication group. It is assumed that if you are using a local certificate to authenticate, it is already added to the FortiGate unit

1. Go to **WAN Opt. & Cache > Authentication Groups**.
2. Select **Create New**.
3. Add a **Name** for the authentication group.

You will select this name when you add the authentication group to a WAN optimization rule.

4. Select the **Authentication Method**.

Select **Certificate** if you want to use a certificate to authenticate and encrypt WAN optimization tunnels. You must select a local certificate that has been added to this FortiGate unit. (To add a local certificate, go to **System > Certificates**.) Other FortiGate units that participate in WAN optimization tunnels with this FortiGate unit must have an authentication group with the same name and certificate.

Select **Pre-shared key** if you want to use a pre-shared key or password to authenticate and encrypt WAN optimization tunnels. You must add the **Password** (or pre-shared key) used by the authentication group. Other FortiGate units that participate in WAN optimization tunnels with this FortiGate unit must have an authentication group with the same name and password. The password must contain at least 6 printable characters and should be known only by network administrators. For optimum protection against currently known attacks, the key should consist of a minimum of 16 randomly chosen alphanumeric characters.

5. Configure **Peer Acceptance** for the authentication group.

Select **Accept Any Peer** if you do not know the peer host IDs or IP addresses of the peers that will use this authentication group. This setting is most often used with FortiGate units that do not have static IP addresses, for example units that use DHCP.

Select **Accept Defined Peers** if you want to authenticate with peers added to the peer list only.

Select **Specify Peer** and select one of the peers added to the peer list to authenticate with the selected peer only.

6. Select **OK**.

7. Add the authentication group to a WAN optimization rule to apply the authentication settings in the authentication group to the rule.

To add an authentication group that uses a certificate- CLI:

Enter the following command to add an authentication group that uses a certificate and can authenticate all peers added to the FortiGate unit configuration.

In this example, the authentication group is named `auth_grp_1` and uses a certificate named `Example_Cert`.

```
config wanopt auth-group
  edit auth_grp_1
    set auth-method cert
    set cert Example_Cert
    set peer-accept defined
  end
```

To add an authentication group that uses a pre-shared key - CLI:

Enter the following command to add an authentication group that uses a pre-shared key and can authenticate only the peer added to the authentication group.

In this example, the authentication group is named `auth_peer`, the peer that the group can authenticate is named `Server_net`, and the authentication group uses `123456` as the pre-shared key. In practice you should use a more secure pre-shared key.

```
config wanopt auth-group
  edit auth_peer
    set auth-method psk
    set psk 123456
    set peer-accept one
    set peer Server_net
  end
```

To add an authentication group that accepts WAN optimization connections from any peer - web-based manager

Add an authentication group that accepts any peer for situations where you do not have the **Peer Host IDs** or **IP Addresses** of the peers that you want to perform WAN optimization with. This setting is most often used for WAN optimization with FortiGate units that do not have static IP addresses, for example units that use DHCP. An authentication group that accepts any peer is less secure than an authentication group that accepts defined peers or a single peer.

The example below sets the authentication method to **Pre-shared key**. You must add the same password to all FortiGate units using this authentication group.

1. Go to **WAN Opt. & Cache > Authentication Groups**.
2. Select **Create New** to add a new authentication group.
3. Configure the authentication group:

Name	Specify any name.
Authentication Method	Pre-shared key
Password	Enter a pre-shared key.
Peer Acceptance	Accept Any Peer

To add an authentication group that accepts WAN optimization connections from any peer - CLI:

In this example, the authentication group is named `auth_grp_1`. It uses a certificate named `WAN_Cert` and accepts any peer.

```
config wanopt auth-group
  edit auth_grp_1
    set auth-method cert
    set cert WAN_Cert
    set peer-accept any
  end
```

Secure tunneling

You can configure WAN optimization rules to use AES-128bit-CBC SSL to encrypt the traffic in the WAN optimization tunnel. WAN optimization uses FortiASIC acceleration to accelerate SSL decryption and encryption of the secure tunnel. Peer-to-peer secure tunnels use the same TCP port as non-secure peer-to-peer tunnels (TCP port 7810).

To use secure tunneling, you must select **Enable Secure Tunnel** in a WAN optimization rule and add an authentication group. The authentication group specifies the certificate or pre-shared key used to set up the secure tunnel. The **Peer Acceptance** setting of the authentication group does not affect secure tunneling.

The FortiGate units at each end of the secure tunnel must have the same authentication group with the same name and the same configuration, including the same pre-shared key or certificate. To use certificates you must install the same certificate on both FortiGate units.

For active-passive WAN optimization you can select **Enable Secure Tunnel** only in the active rule. In peer-to-peer WAN optimization you select **Enable Secure Tunnel** in the WAN optimization rule on both FortiGate units.

For information about active-passive and peer-to-peer WAN optimization, see [Manual \(peer-to-peer\) and active-passive WAN optimization on page 1](#)

For a secure tunneling configuration example, see [Example: Adding secure tunneling to an active-passive WAN optimization configuration on page 1](#).

Monitoring WAN optimization peer performance

The WAN optimization peer monitor lists all of the WAN optimization peers that a FortiGate unit can perform WAN optimization with. These include peers manually added to the configuration as well as discovered peers.

The monitor lists each peer's name, IP address, and peer type. The peer type indicates whether the peer was manually added or discovered. To show WAN optimization performance, for each peer the monitor lists the percent of traffic reduced by the peer in client-side WAN optimization configurations and in server-side configurations (also called gateway configurations).

To view the peer monitor, go to **WAN Opt. & Cache > Peer Monitor**.

Web cache concepts

FortiGate web caching is a form of object caching that accelerates web applications and web servers by reducing bandwidth usage, server load, and perceived latency. Web caching supports caching of HTTP 1.0 and HTTP 1.1 web sites. See [RFC 2616](#) for information about web caching for HTTP 1.1.



Web caching supports caching of Flash content over HTTP but does not cache audio and video streams including Flash videos and streaming content that use native streaming protocols such as RTMP.

The first time a file is received by web caching it is cached in the format it is received in, whether it be compressed or uncompressed. When the same file is requested by a client but in a different compression format, the cached file is converted to the new compressed format before being sent to the client.

There are three significant advantages to using web caching to improve HTTP and WAN performance:

- reduced bandwidth consumption because fewer requests and responses go over the WAN or Internet.
- reduced web server load because there are fewer requests for web servers to handle.
- reduced latency because responses for cached requests are available from a local FortiGate unit instead of from across the WAN or Internet.

You can use web caching to cache any web traffic that passes through the FortiGate unit, including web pages from web servers on a LAN, WAN or on the Internet. You apply web caching by enabling the web caching option in any security policy. When enabled in a security policy, web caching is applied to all HTTP sessions accepted by the security policy. If the security policy is an explicit web proxy security policy, the FortiGate unit caches explicit web proxy sessions.

Turning on web caching for HTTP and HTTPS traffic

Web caching can be applied to any HTTP or HTTPS traffic by enabling web caching in a security policy that accepts the traffic. This includes IPv4, IPv6, WAN optimization and explicit web proxy traffic. Web caching caches all HTTP traffic accepted by a policy on TCP port 80.

You can add web caching to a policy to:

- Cache Internet HTTP traffic for users on an internal network to reduce Internet bandwidth use. Do this by selecting the web cache option for security policies that allow users on the internal network to browse web sites on the Internet.
- Reduce the load on a public facing web server by caching objects on the FortiGate unit. This is a reverse proxy with web caching configuration. Do this by selecting the web cache option for a security policy that allows users on the Internet to connect to the web server.
- Cache outgoing explicit web proxy traffic when the explicit proxy is used to proxy users in an internal network who are connecting to the web servers on the Internet. Do this by selecting the web cache option for explicit web proxy security policies that allow users on the internal network to browse web sites on the Internet.
- Combine web caching with WAN optimization. You can enable web caching in any WAN optimization security policy. This includes manual, active, and passive WAN optimization policies and WAN optimization tunnel policies. You can enable web caching on both the client-side and the server-side FortiGate units or on just one or the other.

For optimum performance you can enable web caching on both the client-side and server-side FortiGate units. In this way only uncached content is transmitted through the WAN optimization tunnel. All cached content is access locally by clients from the client side FortiGate unit.



One important use for web caching is to cache software updates (for example, Windows Updates or iOS updates). When updates occur a large number of users may all be trying to download these updates at the same time. Caching these updates will be a major performance improvement and also have a potentially large impact on reducing Internet bandwidth use. You may want to adjust the maximum cache object size to make sure these updates are cached. See [Turning on web caching for HTTP and HTTPS traffic on page 817](#).

Turning on web caching for HTTPS traffic

Web caching can also cache the content of HTTPS traffic on TCP port 443. With HTTPS web caching, the FortiGate unit receives the HTTPS traffic on behalf of the client, opens up the encrypted traffic and extracts content to be cached. Then FortiGate unit re-encrypts the traffic and sends it on to its intended recipient. It is very similar to a man-in-the-middle attack.

You enable HTTPS web caching from the CLI in a security policy or an explicit proxy policy that accepts the traffic to be cached using `webcache-https`. For a firewall policy:

```
config firewall policy
  edit 0
    .
    .
    .
    set webcache enable
    set webcache-https enable
    .
    .
    .
  end
```

For an explicit web proxy policy:

```
config firewall proxy-policy
  edit 0
    set proxy explicit-web
    .
    .
    .
    set webcache enable
    set webcache-https enable
    .
    .
    .
  end
```



The `webcache-https` field is available only if `webcache` is enabled.

Web caching for HTTPS traffic is not supported if WAN optimization or FTP proxy is enabled: i.e., if `srcintf` is `ftp-proxy` or `wanopt`.

The `any` setting causes the FortiGate unit to re-encrypt the traffic with the FortiGate unit's certificate rather than the original certificate. This configuration can cause errors for HTTPS clients because the name on the certificate does not match the name on the web site.

You can stop these errors from happening by configuring HTTPS web caching to use the web server's certificate by setting `webcache-https` to `ssl-server`. This option is available for both firewall policies and explicit web proxy policies.

```
config firewall policy
  edit 0
    .
    .
    .
    set webcache enable
    set webcache-https enable
    .
    .
    .
  end
```

The `ssl-server` option causes the FortiGate unit to re-encrypt the traffic with a certificate that you imported into the FortiGate unit. You can add certificates using the following command:

```
config firewall ssl-server
  edit corporate-server
    set ip <Web-Server-IP>
    set port 443
    set ssl-mode { full | half}
    set ssl-cert <Web-Server-Cert>
  end
```

Where:

`Web-Server-IP` is the web server's IP address.

`Web-Server-Cert` is a web server certificate imported into the FortiGate unit.

The SSL server configuration also determines whether the SSL server is operating in half or full mode and the port used for the HTTPS traffic.

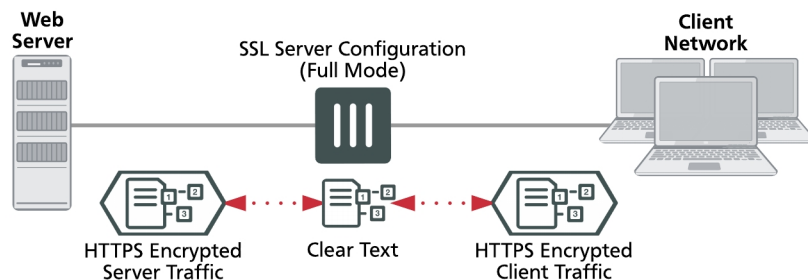
You can add multiple SSL server certificates in this way. When web caching processing an SSL stream if it can find a certificate that matches the web server IP address and port of one of the added SSL servers; that certificate is used to encrypt the SSL traffic before sending it to the client. As a result the client does not generate SSL certificate errors.

Web caching uses the FortiGate unit's FortiASIC to accelerate SSL decryption/encryption performance.

Full mode SSL server configuration

The `ssl-mode` option determines whether the SSL server operates in half or full mode. In full mode the FortiGate unit performs both decryption and encryption of the HTTPS traffic. The full mode sequence is shown below.

Full mode SSL server configuration




In full mode the FortiGate unit is acting as a man in the middle, decrypting and encrypting the traffic. So both the client and the web server see encrypted packets.

Usually the port of the encrypted HTTPS traffic is always 443. However, in the SSL server configuration you can set the port used for HTTPS traffic. This port is not altered by the SSL Server. So for example, if the SSL Server receives HTTPS traffic on port 443, the re-encrypted traffic forwarded to the FortiGate unit to the server or client will still use port 443.

Half mode SSL server configuration

In half mode, the FortiGate unit only performs one encryption or decryption action. If HTTP packets are received, the half mode SSL server encrypts them and converts them to HTTPS packets. If HTTPS packets are received, the SSL server decrypts them and converts them to HTTP packets.

Half mode SSL server configuration



In half mode, the FortiGate unit is acting like an SSL accelerator, offloading HTTPS decryption from the web server to the FortiGate unit. Since FortiGate units can accelerate SSL processing, the end result could be improved web site performance.

Usually the port of the encrypted traffic is always 443. However, in the SSL server configuration you can set the port used for HTTPS traffic. No matter what port is used for the HTTPS traffic, the decrypted HTTP traffic uses port 80.

Changing the ports on which to look for HTTP and HTTPS traffic to cache

By default FortiOS assumes HTTP traffic uses TCP port 80 and HTTPS traffic uses port 443. So web caching caches all HTTP traffic accepted by a policy on TCP port 80 and all HTTPS traffic on TCP port 443. If you want to cache HTTP or HTTPS traffic on other ports, you can enable security profiles for the security policy and configure

a proxy options profile to that looks for HTTP and HTTPS traffic on other TCP ports. To configure a proxy options profile go to **Network > Explicit Proxy**.

Setting the HTTP port to **Any** in a proxy options profile is not compatible with web caching. If you set the HTTP port to any, web caching only caches HTTP traffic on port 80.

Web caching and HA

You can configure web caching on a FortiGate HA cluster. The recommended best practice HA configuration for web caching is active-passive mode. When the cluster is operating, all web caching sessions are processed by the primary unit only. Even if the cluster is operating in active-active mode, HA does not load-balance web caching sessions.

In a cluster, only the primary unit stores the web cache database. The databases is not synchronized to the subordinate units. So, after a failover, the new primary unit must build its web cache.

Web caching and memory usage

To accelerate and optimize disk access and to provide better throughput and less latency, web caching uses provisioned memory to reduce disk I/O and increase disk I/O efficiency. In addition, web caching requires a small amount of additional memory per session for comprehensive flow control logic and efficient traffic forwarding.

When web caching is enabled you will see a reduction in available memory. The reduction increases when more web caching sessions are being processed. If you are thinking of enabling web caching on an operating FortiGate unit, make sure its memory usage is not maxed out during high traffic periods.

In addition to using the system dashboard to see the current memory usage you can use the `get test wad 2` command to see how much memory is currently being used by web caching. See `get test {wad | wccpd} <test_level>` on page 1 for more information.

Changing web cache settings

In most cases, the default settings for the WAN optimization web cache are acceptable. However, you may want to change them to improve performance or optimize the cache for your configuration. To change these settings, go to **WAN Opt. & Cache > Settings**.

From the FortiGate CLI, you can use the `config wanopt webcache` command to change these WAN optimization web cache settings.



For more information about many of these web cache settings, see [RFC 2616](#).

Always revalidate

Select to always revalidate requested cached objects with content on the server before serving them to the client.

Max cache object size

Set the maximum size of objects (files) that are cached. The default size is 512000 KB and the range is 1 to 4294967 KB. This setting determines the maximum object size to store in the web cache. Objects that are larger than this size are still delivered to the client but are not stored in the FortiGate web cache.

For most web traffic the default maximum cache object size is recommended. However, since web caching can also cache larger objects such as Windows updates, Mac OS updates, iOS updates or other updates delivered using HTTP you might want to increase the object size to make sure these updates are cached. Caching these updates can save a lot of Internet bandwidth and improve performance when major updates are released by these vendors.

Negative response duration

Set how long in minutes that the FortiGate unit caches error responses from web servers. If error responses are cached, then subsequent requests to the web cache from users will receive the error responses regardless of the actual object status.

The default is 0, meaning error responses are not cached. The content server might send a client error code (4xx HTTP response) or a server error code (5xx HTTP response) as a response to some requests. If the web cache is configured to cache these negative responses, it returns that response in subsequent requests for that page or image for the specified number of minutes.

Fresh factor

Set the fresh factor as a percentage. The default is 100, and the range is 1 to 100%. For cached objects that do not have an expiry time, the web cache periodically checks the server to see if the objects have expired. The higher the **Fresh Factor** the less often the checks occur.

For example, if you set the **Max TTL** value and **Default TTL** to 7200 minutes (5 days) and set the **Fresh Factor** to 20, the web cache check the cached objects 5 times before they expire, but if you set the **Fresh Factor** to 100, the web cache will check once.

Max TTL

The maximum amount of time (Time to Live) an object can stay in the web cache without the cache checking to see if it has expired on the server. The default is 7200 minutes (120 hours or 5 days) and the range is 1 to 5256000 minutes (5256000 minutes in a year).

Min TTL

The minimum amount of time an object can stay in the web cache before the web cache checks to see if it has expired on the server. The default is 5 minutes and the range is 1 to 5256000 minutes (5256000 minutes in a year).

Default TTL

The default expiry time for objects that do not have an expiry time set by the web server. The default expiry time is 1440 minutes (24 hours) and the range is 1 to 5256000 minutes (5256000 minutes in a year).

Proxy FQDN

The fully qualified domain name (FQDN) for the proxy server. This is the domain name to enter into browsers to access the proxy server. This field is for information only can be changed from the explicit web proxy configuration.

Max HTTP request length

The maximum length of an HTTP request that can be cached. Larger requests will be rejected. This field is for information only can be changed from the explicit web proxy configuration.

Max HTTP message length

The maximum length of an HTTP message that can be cached. Larger messages will be rejected. This field is for information only can be changed from the explicit web proxy configuration.

Ignore

Select the following options to ignore some web caching features.

If-modified-since	By default, if the time specified by the if-modified-since (IMS) header in the client's conditional request is greater than the last modified time of the object in the cache, it is a strong indication that the copy in the cache is stale. If so, HTTP does a conditional GET to the Overlay Caching Scheme (OCS), based on the last modified time of the cached object. Enable ignoring if-modified-since to override this behavior.
HTTP 1.1 conditionals	HTTP 1.1 provides additional controls to the client over the behavior of caches toward stale objects. Depending on various cache-control headers, the FortiGate unit can be forced to consult the OCS before serving the object from the cache. For more information about the behavior of cache-control header values, see RFC 2616 . Enable ignoring HTTP 1.1 Conditionals to override this behavior.
Pragma-no-cache	Typically, if a client sends an HTTP GET request with a pragma no-cache (PNC) or cache-control no-cache header, a cache must consult the OCS before serving the content. This means that the FortiGate unit always re-fetches the entire object from the OCS, even if the cached copy of the object is fresh. Because of this behavior, PNC requests can degrade performance and increase server-side bandwidth utilization. However, if you enable ignoring Pragma-no-cache, then the PNC header from the client request is ignored. The FortiGate unit treats the request as if the PNC header is not present.
IE Reload	Some versions of Internet Explorer issue Accept / header instead of Pragma no-cache header when you select Refresh . When an Accept header has only the / value, the FortiGate unit treats it as a PNC header if it is a type-N object. Enable ignoring IE reload to cause the FortiGate unit to ignore the PNC interpretation of the Accept / header.

Cache expired objects

Applies only to type-1 objects. When this option is selected, expired type-1 objects are cached (if all other conditions make the object cacheable).

Revalidated pragma-no-cache

The pragma-no-cache (PNC) header in a client's request can affect how efficiently the FortiGate unit uses bandwidth. If you do not want to completely ignore PNC in client requests (which you can do by selecting to ignore Pragma-no-cache, above), you can nonetheless lower the impact on bandwidth usage by selecting **Revalidate Pragma-no-cache**.

When you select **Revalidate Pragma-no-cache**, a client's non-conditional PNC-GET request results in a conditional GET request sent to the OCS if the object is already in the cache. This gives the OCS a chance to return the 304 Not Modified response, which consumes less server-side bandwidth, because the OCS has not been forced to otherwise return full content.

By default, **Revalidate Pragma-no-cache** is disabled and is not affected by changes in the top-level profile.

Most download managers make byte-range requests with a PNC header. To serve such requests from the cache, you should also configure byte-range support when you configure the **Revalidate pragma-no-cache** option.

Web cache configuration

Forwarding URLs to forwarding servers and exempting web sites from web caching

You can go to **Network > Explicit Proxy** and use the URL match list to forward URL patterns to forwarding servers and create a list of URLs that are exempt from web caching.

Forwarding URLs and URL patterns to forwarding servers

As part of configuring the explicit web proxy you can configure proxy chaining by adding web proxy forwarding servers. See [Proxy chaining \(web proxy forwarding servers\)](#).

You can then use the URL match list to always forward explicit web proxy traffic destined for configured URLs or URL patterns to one of these forwarding servers. For example, you might want to forward all traffic for a specific country to a proxy server located in that country.

To forward traffic destined for a URL to a forwarding server that you have already added, go to **Network > Explicit Proxy** and select **Create New**. Add a name for the URL match entry and enter the URL or URL pattern. You can use wildcards such as * and ? and you can use a numeric IP address. Select **Forward to Server** and select a web proxy forwarding server from the list.

You can also exempt the URL or URL pattern from web caching.

Use the following command to forward all .ca traffic to a proxy server and all .com traffic to another proxy server.

```
config web-proxy url-match
  edit "com"
    set forward-server "server-commercial"
    set url-pattern "com"
  next
  edit "ca"
    set forward-server "server-canada"
    set url-pattern "ca"
  next
  edit "www.google.ca"
    set cache-exemption enable
    set url-pattern "www.google.ca"
  next
end
```

Exempting web sites from web caching

You may want to exempt some URLs from web caching for a number of reasons. For example, if your users access websites that are not compatible with FortiGate web caching you can add the URLs of these web sites to the web caching exempt list. You can add URLs and numeric IP addresses to the web cache exempt list.

You can also add URLs to the web cache exempt list by going to **Network > Explicit Proxy**, going to the **URL Match List**

URL Match List

+ Create New

Edit

Delete

Name	URL Pattern	Cache Exemption	Forward Server	Status	Comments
No matching entries found					

and selecting **Create New**. Add a URL pattern to be exempt and select **Exempt from Cache**.

New URL Match Entry

Name

Comments

Comments 0/255

URL Pattern

Forward to Server

☐

Exempt from Cache

☒

Enable this URL

☒

You can also add URLs and addresses to be exempt from caching using the CLI. Enter the following command to add `www.example.com` to the web cache exempt list:

```
config web-proxy url-match
  set cache-exemption enable
  set url-pattern www.example.com
end
```

Exempting specific files from caching

You can exempt files from being cached, so long as you specify its full URL. Enter the following command to add the URL, with the file extension (in this example, `.exe`), to the web cache exempt list:

```
config web-proxy url-match
  edit "exe"
    set url-pattern "iavs9x.u.avast.com/custom/iavs9x/20160613t1237z/avast_free_
      antivirus_setup_online.exe"
    set cache-exemption enable
  next
end
```



You cannot use wildcards to exempt file extensions in general from caching.

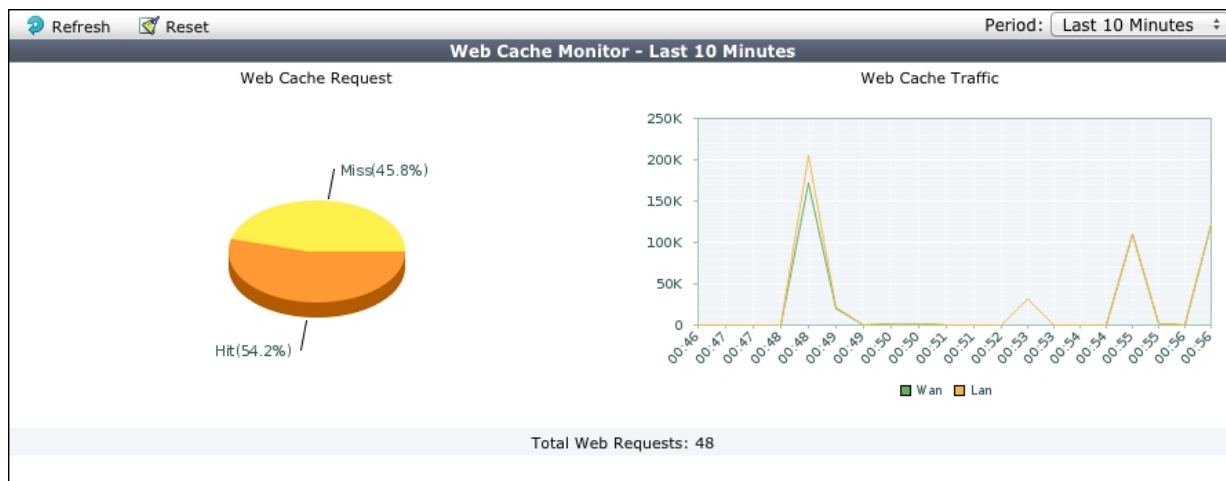
Monitoring web caching performance

The web cache monitor shows the percentage of web cache requests that retrieved content from the cache (hits) and the percentage that did not receive content from the cache (misses). A higher the number of hits usually indicates that the web cache is being more effective at reducing WAN traffic.

The web cache monitor also shows a graph of web traffic on the WAN and LAN. A lower WAN line on the graph indicates the web cache is reducing traffic on the WAN. The web cache monitor also displays the total number of web requests processed by the web cache.

To view the web cache monitor, go to **Monitor > Cache Monitor**.

Web cache monitor



Example web caching of HTTP and HTTPS Internet content for users on an internal network

This example describes how to configure web caching of HTTP and HTTPS for users on a private network connecting to the Internet.

Network topology and assumptions

This example includes a client network with subnet address 10.31.101.0 connecting to web servers on the Internet. All of the users on the private network access the Internet through a single general security policy on the FortiGate unit that accepts all sessions connecting to the Internet. Web caching for HTTP and HTTPS traffic is added to this security policy.


Since users on the private network have unrestricted access to the Internet and can be accessing many web servers the `webcache-https` is set to `any` and users may see error messages on their web browsers when accessing HTTPS content.

The GUI is less versatile than the CLI so the example instructions for the GUI give settings for one port for each protocol, while the CLI example shows how to use multiple ports.

The example also describes how to configure the security policy to cache HTTP traffic on port 80 and 8080 in the CLI, by adding a proxy options profile that looks for HTTP traffic on TCP ports 80 and 8080. The example also

describes how to configure the security policy to cache HTTPS traffic on port 443 and 8443 using the same proxy options profile.

Example web caching topology



General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Add HTTP web caching to the security policy that all users on the private network use to connect to the Internet.
2. Add HTTPS web caching.
3. Add a protocol options profile to look for HTTP traffic on ports 80 and 8080 and HTTPS traffic on ports 443 and 8443 and add this protocol options profile to the security policy.

If you perform any additional actions between procedures, your configuration may have different results.

Configuration steps - web-based manager

Use the following steps to configure the example configuration from the FortiGate web-based manager.

To add HTTP web caching to a security policy

1. Go to **Policy & Objects > IPv4 Policy** and add a security policy that allows all users on the internal network to access the Internet.

Incoming Interface	Internal
Outgoing Interface	wan1
Source	all
Destination	all
Schedule	always
Service	ALL
Action	ACCEPT

2. Toggle **NAT** to enabled, and select **Use Outgoing Interface Address**.
3. Turn on **Web cache**.
4. Select **OK**.

To add HTTPS web caching

1. From the CLI enter the following command to add HTTPS web caching to the policy.

Assume the index number of the policy is 5.

```
config firewall policy
edit 5
set webcache-https any
end
```

To cache HTTP traffic on port 80 and HTTPS on 8443

1. Go to **Network > Explicit Proxy** and edit the Explicit Proxy options profile.
2. Under **Explicit Web Proxy**,
 - For the **HTTP port**, enter 80.
 - For **HTTPS port**, select **Specify** and enter 8443 in the field.
3. Click on **Apply**.



You need to use the CLI to add the protocol options profile unless you also add a security profile that uses proxy-based inspection.

Configuration steps - CLI

Use the following steps to configure the example configuration from the FortiGate CLI.

To add HTTP and HTTPS web caching to a security policy

1. Enter the following command to add a security policy that allows all users on the internal network to access the Internet and that includes web caching of HTTP and HTTPS traffic.

```
config firewall policy
edit 0
set srcintf internal
set srcaddr all
set dstintf wan1
set dstintf all
set schedule always
set service ANY
set action accept
set nat enable
set webcache enable
set webcache-https any
end
```

To cache HTTP traffic on port 80 and 8080 and HTTPS traffic on ports 443 and 8443

1. Enter the following command to edit the **default** proxy options profile to configure it to look for HTTP traffic on ports 80 and 8080:

```
config firewall profile-protocol-options
edit default
config http
set status enable
set ports 80 8080
```

```
end
```

2. Enter the following command to edit the **certificate-inspection** SSL SSH options profile to configure it to look for HTTPS traffic on ports 443 and 8443:

```
config firewall ssl-ssh-profile
edit certificate-inspection
config https
set status certificate-inspection
set ports 443 8443
end
```

3. Enter the following command to add the **default** proxy options profile and the **certificate-inspection** SSL SSH profile to the firewall policy.

```
config firewall policy
edit 5
set utm-status enable
set profile-protocol-options default
set ssl-ssh-profile certificate-inspection
end
```

Example reverse proxy web caching and SSL offloading for an Internet web server using a static one-to-one virtual IP

This section describes configuring SSL offloading for a reverse proxy web caching configuration using a static one-to-one firewall virtual IP (VIP). While the static one-to-one configuration described in this example is valid, its also common to change the destination port of the unencrypted HTTPS traffic to a commonly used HTTP port such as 8080 using a port forwarding virtual IP.

Network topology and assumptions

In this configuration, clients on the Internet use HTTP and HTTPS to browse to a web server that is behind a FortiGate unit. A policy added to the FortiGate unit forwards the HTTP traffic to the web server. The policy also offloads HTTPS decryption and encryption from the web server so the web server only sees HTTP traffic.

The FortiGate unit also caches HTTP and HTTPS pages from the web server so when users access cached pages the web server does not see the traffic. Replies to HTTPS sessions are encrypted by the FortiGate unit before returning to the clients.

In this configuration, the FortiGate unit is operating as a web cache in reverse proxy mode. Reverse proxy caches can be placed directly in front of a web server. Web caching on the FortiGate unit reduces the number of requests that the web server must handle, therefore leaving it free to process new requests that it has not serviced before.

Using a reverse proxy configuration:

- avoids the capital expense of additional web servers by increasing the capacity of existing servers
- serves more requests for static content from web servers
- serves more requests for dynamic content from web servers
- reduces operating expenses including the cost of bandwidth required to serve content
- accelerates the response time of web servers and of page download times to end users.

When planning a reverse proxy implementation, the web server's content should be written so that it is "cache aware" to take full advantage of the reverse proxy cache.

In reverse proxy mode, the FortiGate unit functions more like a web server for clients on the Internet. Replicated content is delivered from the proxy cache to the external client without exposing the web server or the private network residing safely behind the firewall.

In this example, the site URL translates to IP address 192.168.10.1, which is the port2 IP address of the FortiGate unit. The port2 interface is connected to the Internet.

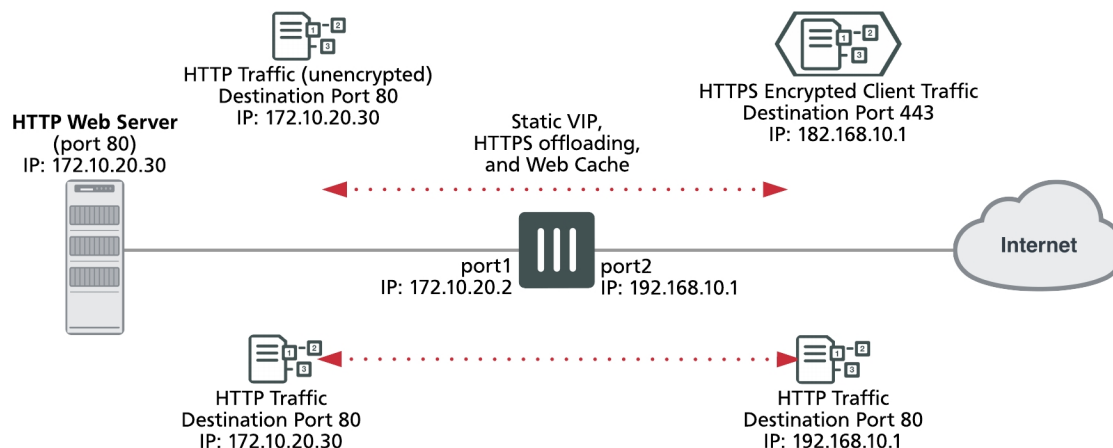
This example assumes that all HTTP traffic uses port 80 and all HTTPS traffic uses port 443.

The FortiGate unit includes the web server CA and an SSL server configuration for IP address 172.10.20.30 and port to 443. The name of the file containing the CA is Rev_Proxy_Cert_1.crt.

The destination address of incoming HTTP and HTTPS sessions is translated to the IP address of the web server using a static one-to-one virtual IP that performs destination address translation (DNAT) for the HTTP packets. The DNAT translates the destination address of the packets from 192.168.10.1 to 172.10.20.30 but does not change the destination port number.

When the SSL server on the FortiGate unit decrypts the HTTPS packets their destination port is changed to port 80.

Reverse proxy web caching and SSL offloading for an Internet web server using static one-to-one virtual IPs



General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Configure the FortiGate unit as a reverse proxy web cache server.
2. Configure the FortiGate unit for SSL offloading of HTTPS traffic.
3. Add an SSL server to offload SSL encryption and decryption for the web server.

Also note that if you perform any additional actions between procedures, your configuration may have different results.

Configuration steps - web-based manager

To configure the FortiGate unit as a reverse proxy web cache server

1. Go to **Policy & Objects > Virtual IPs** and select **Create New** to add a static NAT virtual IP that translates destination IP addresses from 192.168.10.1 to 172.10.20.30 (and does not translate destination ports):

VIP Type	IPv4
Name	Reverse_proxy_VIP
Interface	port2
Type	Static NAT
Optional Filters	Do not select.
External IP Address/Range	192.168.10.1
Mapped IP Address/Range	172.10.20.30
Port Forwarding	Do not select.

2. Select **OK**.
3. Go to **Policy & Objects > IPv4 Policy** and select **Create New** to add a port2 to port1 security policy that accepts HTTP and HTTPS traffic from the Internet.

Do not select security profiles. Set the destination address to the virtual IP. You do not have to enable NAT.

Incoming Interface	port2
Outgoing Interface	port1
Source	all
Destination	Reverse_proxy_VIP
Schedule	always
Service	HTTP HTTPS
Action	ACCEPT

4. Turn on **Web Cache**.
5. Select **OK**.
6. From the CLI enter the following command to add HTTPS web caching to the security policy

Assume the index number of the policy is 5.

```
config firewall policy
edit 5
set webcache-https ssl-server
end
```

To configure the FortiGate unit to offload SSL encryption and cache HTTPS content

1. Go to **System > Certificates** and select **Import** to import the web server's CA.

For **Type**, select **Local Certificate**. Select the **Browse** button to locate the file (example file name: Rev_Proxy_Cert_1.crt).

The certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported.

2. Select **OK** to import the certificate.
3. From the CLI, enter the following command to add the SSL server and to add the server's certificate to the SSL server.

The SSL server `ip` must match the destination address of the SSL traffic after being translated by the virtual IP (172.10.20.30) and the SSL server `port` must match the destination port of the SSL traffic (443). The SSL server operates in half mode since it performs a single-step conversion (HTTPS to HTTP or HTTP to HTTPS).

```
config firewall ssl-server
  edit rev_proxy_server
    set ip 172.10.20.30
    set port 443
    set ssl-mode half
    set ssl-cert Rev_Proxy_Cert_1
  end
```

Configuration steps - CLI

To configure the FortiGate unit as a reverse proxy web cache server

1. Enter the following command to add a static NAT virtual IP that translates destination IP addresses from 192.168.10.1 to 172.10.20.30 (and does not translate destination ports):

```
config firewall vip
  edit Reverse_proxy_VIP
    set extintf port2
    set type static-nat
    set extip 192.168.10.1
    set mappedip 172.10.20.30
  end
```

2. Enter the following command to add a port2 to port1 security policy that accepts HTTP and HTTPS traffic from the Internet. Enable web caching and HTTPS web caching.

Do not select security profiles. Set the destination address to the virtual IP. You do not have to enable NAT.

```
config firewall policy
  edit 0
    set srcintf port2
    set srcaddr all
    set dstintf port1
    set dstaddr Reverse_proxy_VIP
    set schedule always
    set service HTTP HTTPS
    set action accept
    set webcache enable
    set webcache-https ssl-server
  end
```


To add an SSL server to offload SSL encryption and decryption for the web server

1. Place a copy of the web server's CA (file name `Rev_Proxy_Cert_1.crt`) in the root folder of a TFTP server.
2. Enter the following command to import the web server's CA from a TFTP server. The IP address of the TFTP server is 10.31.101.30:

```
execute vpn certificate local import tftp Rev_Proxy_Cert_1.crt 10.31.101.30
```

The certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported.

3. From the CLI, enter the following command to add the SSL server.

The SSL server `ip` must match the destination address of the SSL traffic after being translated by the virtual IP (172.10.20.30) and the SSL server `port` must match the destination port of the SSL traffic (443). The SSL server operates in half mode since it performs a single-step conversion (HTTPS to HTTP or HTTP to HTTPS).

```
config firewall ssl-server
edit rev_proxy_server
set ip 172.10.20.30
set port 443
set ssl-mode half
set ssl-cert Rev_Proxy_Cert_1
end
```

4. Configure other `ssl-server` settings that you may require for your configuration.

Using a FortiCache as a cache service

Some FortiGate devices don't have sufficient memory or disk space to run a cache service. This feature allows a FortiGate to connect to a FortiCache that has a higher cache capability than most FortiGates.

Syntax:

```
config wanopt remote-storage
set status {enable|disable}
set local-cache-id <name ID for connection>
set remote-cache-id <ID of the remote device>
set remote-cache-ip <IP address of the remote device>
end
```

Option	Description
status	Enable or disable whether the FortiGate uses a remote caching device as web-cache storage. If disabled, uses local disk(s) as web storage.
local-cache-id	ID that this device uses to connect to the remote caching device
remote-cache-id	ID of the remote caching device that this FortiGate connects to
remote-cache-ip	IP address of the remote caching device that this FortiGate connects to.

WCCP concepts

The Web Cache Communication Protocol (WCCP) can be used to provide web caching with load balancing and fault tolerance. In a WCCP configuration, a WCCP server receives HTTP requests from user's web browsers and redirects the requests to one or more WCCP clients. The clients either return cached content or request new content from the destination web servers before caching it and returning it to the server which in turn returns the content to the original requestor. If a WCCP configuration includes multiple WCCP clients, the WCCP server load balances traffic among the clients and can detect when a client fails and failover sessions to still operating clients. WCCP is described by the [Web Cache Communication Protocol Internet draft](#).

The sessions that are cached by WCCP depend on the configuration of the WCCP clients. If the client is a FortiGate unit, you can configure the port numbers and protocol number of the sessions to be cached. For example, to cache HTTPS traffic on port 443 the WCCP client port must be set to 443 and protocol must be set to 6. If the WCCP client should also cache HTTPS traffic on port 993 the client ports option should include both port 443 and 993.

On a FortiGate unit, WCCP sessions are accepted by a security policy before being cached. If the security policy that accepts sessions that do not match the port and protocol settings in the WCCP clients the traffic is dropped.

WCCP is configured per-VDOM. A single VDOM can operate as a WCCP server or client (not both at the same time). FortiGate units are compatible with third-party WCCP clients and servers. If a FortiGate unit is operating as an Internet firewall for a private network, you can configure it to cache and serve some or all of the web traffic on the private network using WCCP by adding one or more WCCP clients, configuring WCCP server settings on the FortiGate unit and adding WCCP security policies that accept HTTP session from the private network.

FortiGate units support WCCPv1 and WCCPv2. A FortiGate unit in NAT/Route or transparent mode can operate as a WCCP server. To operate as a WCCP client a FortiGate unit must be in NAT/Route mode. FortiGate units communicate between WCCP servers and clients over UDP port 2048. This communication can be encapsulated in a GRE tunnel or just use layer 2 forwarding.



A WCCP server can also be called a WCCP router. A WCCP client can also be called a WCCP cache engine.

WCCP Cisco to FortiGate client using L2-forwarding tunneling

FortiGate supports the option of using Mask mode, in addition to Hash mode, when operating as a WCCP client using L2 forwarding. As a result, you can configure a WCCP FortiGate client to connect to a Cisco Nexus, which doesn't accept the Hash mode assignment method, using the Mask mode assignment method.

WCCP configuration

WCCP configuration overview

To configure WCCP you must create a service group that includes WCCP servers and clients. WCCP servers intercept sessions to be cached (for example, sessions from users browsing the web from a private network). To intercept sessions to be cached the WCCP server must include a security policy that accepts sessions to be cached and WCCP must be enabled in this security policy.

The server must have an interface configured for WCCP communication with WCCP clients. That interface sends and receives encapsulated GRE traffic to and from WCCP clients. The server must also include a WCCP service group that includes a service ID and the addresses of the WCCP clients as well as other WCCP configuration options.

To use a FortiGate unit as a WCCP client, the FortiGate unit must be set to be a WCCP client (or cache engine). You must also configure an interface on the client for WCCP communication. The client sends and receives encapsulated GRE traffic to and from the WCCP server using this interface.

The client must also include a WCCP service group with a service ID that matches a service ID on the server. The client service group also includes the IP address of the servers in the service group and specifies the port numbers and protocol number of the sessions that will be cached on the client.

When the client receives sessions from the server on its WCCP interface, it either returns cached content over the WCCP interface or connects to the destination web servers using the appropriate interface depending on the client routing configuration. Content received from web servers is then cached by the client and returned to the WCCP server over the WCCP link. The server then returns the received content to the initial requesting user web browser.

Finally you may also need to configure routing on the server and client FortiGate units and additional security policies may have to be added to the server to accept sessions not cached by WCCP.

WCCP service groups, service numbers, service IDs and well known services

A FortiGate unit configured as a WCCP server or client can include multiple server or client configurations. Each of these configurations is called a WCCP service group. A service group consists of one or more WCCP servers (or routers) and one or more WCCP clients working together to cache a specific type of traffic. The service group configuration includes information about the type of traffic to be cached, the addresses of the WCCP clients and servers and other information about the service.

A service group is identified with a numeric WCCP service ID (or service number) in the range 0 to 255. All of the servers and clients in the same WCCP service group must have service group configurations with the same WCCP service ID.

The value of the service ID provides some information about the type of traffic to be cached by the service group. Service IDs in the range 0 to 50 are reserved for well known services. A well known service is any service that is defined by the WCCP standard as being well known. Since the service is well known, just the service ID is required to identify the traffic to be cached.

Even though the well known service ID range is 0 to 50, at this time only one well known service has been defined. Its service ID 0, which is used for caching HTTP (web) traffic.

So to configure WCCP to cache HTTP sessions you can add a service group to the WCCP router and WCCP clients with a service ID of 0. No other information about the type of traffic to cache needs to be added to the service group.

Since service IDs 1 to 50 are reserved for well know services and since these services are not defined yet, you should not add service groups with IDs in the range 1 to 50.



FortiOS does allow you to add service groups with IDs between 1 and 50. Since these service groups have not been assigned well known services, however, they will not cache any sessions. Service groups with IDs 51 to 255 allow you to set the port numbers and protocol number of the traffic to be cached. So you can use service groups with IDs 51 to 255 to cache different kinds of traffic based on port numbers and protocol number of the traffic. Service groups 1 to 50; however, do not allow you to set port numbers or protocol numbers so cannot be used to cache any traffic.

To cache traffic other than HTTP traffic you must add service groups with IDs in the range 51 to 255. These service group configurations must include the port numbers and protocol number of the traffic to be cached. It is the port and protocol number configuration in the service group that determines what traffic will be cached by WCCP.

Example WCCP server and client configuration for caching HTTP sessions (service ID = 0)

Enter the following command to add a WCCP service group to a WCCP server that caches HTTP sessions. The IP address of the server is 10.31.101.100 and the WCCP clients are on the 10.31.101.0 subnet. The service

ID of this service group is 0.

```
config system wccp
  edit 0
    set router-id 10.31.101.100
    set server-list 10.31.101.0 255.255.255.0
  end
```

Enter the following commands to configure a FortiGate unit to operate as a WCCP client and add a service group that configures the client to cache HTTP sessions. The IP address of the server is 10.31.101.100 and IP address of this WCCP clients is 10.31.101.1 subnet. The service ID of this service group is 0.

```
config system settings
  set wccp-cache-engine enable
end

config system wccp
  edit 0
    set cache-id 10.31.101.1
    set router-list 10.31.101.100
  end
```



You cannot enter the `wccp-cache-engine enable` command if you have already added a WCCP service group. When you enter this command an interface named `w.<vdom_name>` is added to the FortiGate configuration (for example `w.root`). All traffic redirected from a WCCP router is considered to be received at this interface of the FortiGate unit operating as a WCCP client. A default route to this interface with lowest priority is added.

Example WCCP server and client configuration for caching HTTPS sessions

Enter the following command to add a service group to a WCCP server that caches HTTPS content on port 443 and protocol 6. The IP address of the server is 10.31.101.100 and the WCCP clients are on the 10.31.101.0 subnet. The service ID of this service group is 80.

```
config system settings
    set wccp-cache-engine enable
end

config system wccp
    edit 80
        set router-id 10.31.101.100
        set server-list 10.31.101.0 255.255.255.0
        set ports 443
        set protocol 6
    end
```

Enter the following commands to configure a FortiGate unit to operate as a WCCP client and add a service group that configures client to cache HTTPS sessions on port 443 and protocol 6. The IP address of the server is 10.31.101.100 and IP address of this WCCP clients is 10.31.101.1 subnet. The service ID of this service group must be 80 to match the service ID added to the server.

```
config system settings
    set wccp-cache-engine enable
end

config system wccp
    edit 80
        set cache-id 10.31.101.1
        set router-list 10.31.101.100
        set ports 443
        set protocol 6
    end
```

Example WCCP server and client configuration for caching HTTP and HTTPS sessions

You could do this by configuring two WCCP service groups as described in the previous examples. Or you could use the following commands to configure one service group for both types of traffic. The example also caches HTTP sessions on port 8080.

Enter the following command to add a service group to a WCCP server that caches HTTP sessions on ports 80 and 8080 and HTTPS sessions on port 443. Both of these protocols use protocol number 6. The IP address of the server is 10.31.101.100 and the WCCP clients are on the 10.31.101.0 subnet. The service ID of this service group is 90.

```
config system wccp
    edit 90
        set router-id 10.31.101.100
        set server-list 10.31.101.0 255.255.255.0
        set ports 443 80 8080
        set protocol 6
    end
```

Enter the following commands to configure a FortiGate unit to operate as a WCCP client and add a service group that configures client to cache HTTP sessions on port 80 and 8080 and HTTPS sessions on port 443. The IP

address of the server is 10.31.101.100 and IP address of this WCCP clients is 10.31.101.1 subnet. The service ID of this service group must be 90 to match the service ID added to the server.

```
config system settings
    set wccp-cache-engine enable
end
config system wccp
    edit 90
        set cache-id 10.31.101.1
        set router-list 10.31.101.100
        set ports 443 80 8080
        set protocol 6
    end
```

Other WCCP service group options

In addition to using WCCP service groups to define the types of traffic to be cached by WCCP the following options are available for servers and clients.

Server configuration options

The server configuration must include the `router-id`, which is the WCCP server IP address. This is the IP address of the interface that the server uses to communicate with WCCP clients.

The `group-address` is used for multicast WCCP configurations to specify the multicast addresses of the clients.

The `server-list` defines the IP addresses of the WCCP clients that the server can connect to. Often the server list can be the address of the subnet that contains the WCCP clients.

The `authentication` option enables or disables authentication for the WCCP service group. Authentication must be enabled on all servers and clients in a service group and members of the group must have the same password.

The `forward-method` option specifies the protocol used for communication between the server and clients. The default forwarding method is GRE encapsulation. If required by your network you can also select to use unencapsulated layer-2 packets instead of GRE or select any to allow both. The `return-method` allows you to specify the communication method from the client to the server. Both GRE and layer-2 are supported.

The `assignment-method` determines how the server load balances sessions to the clients if there are multiple clients. Load balancing can be done using hashing or masking.

Client configuration options

The client configuration includes the `cache-id` which is the IP address of the FortiGate interface of the client that communicates with WCCP server. The `router-list` option is the list of IP addresses of the WCCP servers in the WCCP service group.

The `ports` option lists the port numbers of the sessions to be cached by the client and the `protocol` sets the protocol number of the sessions to be cached. For TCP sessions the protocol is 6.

The `service-type` option can be auto, dynamic or standard. Usually you would not change this setting.

The client configuration also includes options to influence load balancing including the `primary-hash`, `priority`, `assignment-weight` and `assignment-bucket-format`.

Example caching HTTP sessions on port 80 using WCCP

In this example configuration (shown below), a FortiGate unit with host name WCCP_srv is operating as an Internet firewall for a private network is also configured as a WCCP server. The port1 interface of WCCP_srv is connected to the Internet and the port2 interface is connected to the internal network.

All HTTP traffic on port 80 that is received at the port2 interface of WCCP_srv is accepted by a port2 to port1 security policy with WCCP enabled. All other traffic received at the port2 interface is allowed to connect to the Internet by adding a general port2 to port1 security policy below the HTTP on port 80 security policy.

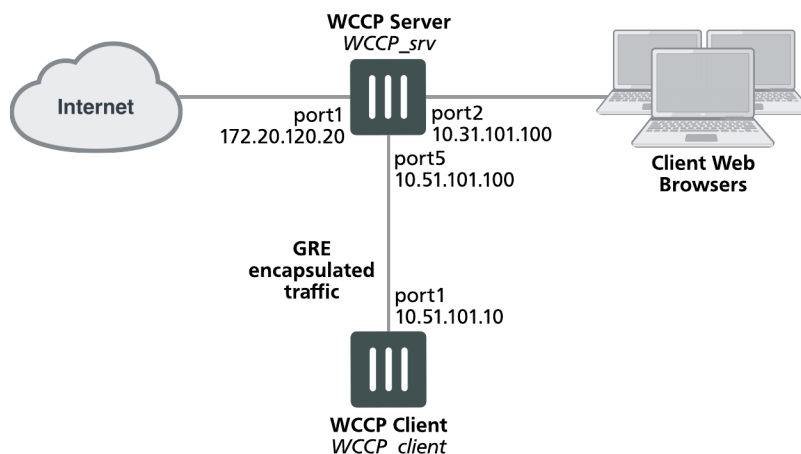
A WCCP service group is added to WCCP_srv with a service ID of 0 for caching HTTP traffic on port 80. The port5 interface of WCCP_srv is configured for WCCP communication.

A second FortiGate unit with host name WCCP_client is operating as a WCCP client. The port1 interface of WCCP_client is connected to port5 of WCCP_srv and is configured for WCCP communication.

WCCP_client is configured to cache HTTP traffic because it also has a WCCP service group with a service ID of 0.

WCCP_client connects to the Internet through WCCP_srv. To allow this, a port5 to port1 security policy is added to WCCP_srv.

FortiGate WCCP server and client configuration



Configuring the WCCP server (WCCP_srv)

Use the following steps to configure WCCP_srv as the WCCP server for the example network. The example steps only describe the WCCP-related configuration.

To configure WCCP_srv as a WCCP server

1. Add a port2 to port1 security policy that accepts HTTP traffic on port 80 and is configured for WCCP:

```

config firewall policy
  edit 0
    set srtintf port2
    set dstintf port1
    set srcaddr all
    set dstaddr all
  
```

```

        set action accept
        set schedule always
        set service HTTP
        set wccp enable
        set nat enable
    end

```

2. Add another port2 to port1 security policy to allow all other traffic to connect to the Internet.

```

config firewall policy
    edit 0
        set srtintf port2
        set dstintf port1
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule always
        set service ANY
        set nat enable
    end

```

3. Move this policy below the WCCP policy in the port2 to port1 policy list.

4. Enable WCCP on the port5 interface.

```

config system interface
    edit port5
        set wccp enable
    end

```

5. Add a WCCP service group with service ID 0.

```

config system wccp
    edit 0
        set router-id 10.51.101.100
        set server-list 10.51.101.0 255.255.255.0
    end

```

6. Add a firewall address and security policy to allow the WCCP_client to connect to the internet.

```

config firewall address
    edit WCCP_client_addr
        set subnet 10.51.101.10
    end
config firewall policy
    edit 0
        set srtintf port5
        set dstintf port1
        set srcaddr WCCP_client_addr
        set dstaddr all
        set action accept
        set schedule always
        set service ANY
        set nat enable
    end

```

Configuring the WCCP client (WCCP_client)

Use the following steps to configure WCCP_client as the WCCP client for the example network. The example steps only describe the WCCP-related configuration.

To configure WCCP_client as a WCCP client

1. Configure WCCP_client to operate as a WCCP client.

```
config system settings
    set wccp-cache-engine enable
end
```



You cannot enter the `wccp-cache-engine enable` command if you have already added a WCCP service group. When you enter this command an interface named `w.<vdom_name>` is added to the FortiGate configuration (for example `w.root`). All traffic redirected from a WCCP router is considered to be received at this interface of the FortiGate unit operating as a WCCP client. A default route to this interface with lowest priority is added.

2. Enable WCCP on the port1 interface.

```
config system interface
    edit port1
        set wccp enable
    end
```

3. Add a WCCP service group with service ID 0.

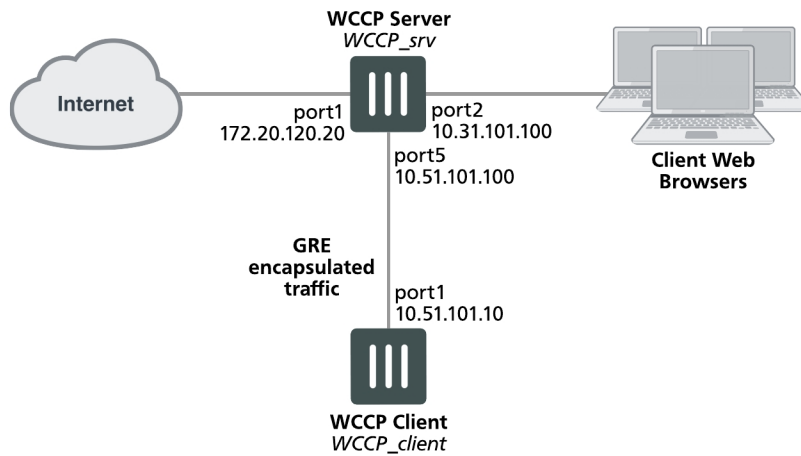
```
config system wccp
    edit 0
        set cache-id 10.51.101.10
        set router-list 10.51.101.100
    end
```

Example caching HTTP sessions on port 80 and HTTPS sessions on port 443 using WCCP

This example configuration is the same as that described in [Example caching HTTP sessions on port 80 and HTTPS sessions on port 443 using WCCP on page 842](#) except that WCCP now also cached HTTPS traffic on port 443. To cache HTTP and HTTPS traffic the WCCP service group must have a service ID in the range 51 to 255 and you must specify port 80 and 443 and protocol 6 in the service group configuration of the WCCP client.

Also the security policy on the WCCP_srv that accepts sessions from the internal network to be cached must accept HTTP and HTTPS sessions.

FortiGate WCCP server and client configuration



Configuring the WCCP server (WCCP_srv)

Use the following steps to configure WCCP_srv as the WCCP server for the example network. The example steps only describe the WCCP-related configuration.

To configure WCCP_srv as a WCCP server

1. Add a port2 to port1 security policy that accepts HTTP traffic on port 80 and HTTPS traffic on port 443 and is configured for WCCP:

```

config firewall policy
  edit 0
    set srtintf port2
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service HTTP HTTPS
    set wccp enable
    set nat enable
  end

```

2. Add another port2 to port1 security policy to allow all other traffic to connect to the Internet.

```

config firewall policy
  edit 0
    set srtintf port2
    set dstintf port1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY

    set nat enable
  end

```

3. Move this policy below the WCCP policy in the port2 to port1 policy list.
4. Enable WCCP on the port5 interface.

```
config system interface
edit port5
set wccp enable
end
```

5. Add a WCCP service group with service ID 90 (can be any number between 51 and 255).

```
config system wccp
edit 90
set router-id 10.51.101.100
set server-list 10.51.101.0 255.255.255.0
end
```

6. Add a firewall address and security policy to allow the WCCP_client to connect to the internet.

```
config firewall address
edit WCCP_client_addr
set subnet 10.51.101.10
end
config firewall policy
edit 0
set srtintf port5
set dstintf port1
set srcaddr WCCP_client_addr
set dstaddr all
set action accept
set schedule always
set service ANY
set nat enable
end
```

Configuring the WCCP client (WCCP_client)

Use the following steps to configure WCCP_client as the WCCP client for the example network. The example steps only describe the WCCP-related configuration.

To configure WCCP_client as a WCCP client

1. Configure WCCP_client to operate as a WCCP client.

```
config system settings
set wccp-cache-engine enable
end
```



You cannot enter the `wccp-cache-engine enable` command if you have already added a WCCP service group. When you enter this command an interface named `w.<vdom_name>` is added to the FortiGate configuration (for example `w.root`). All traffic redirected from a WCCP router is considered to be received at this interface of the FortiGate unit operating as a WCCP client. A default route to this interface with lowest priority is added.

2. Enable WCCP on the port1 interface.

```
config system interface
edit port1
set wccp enable
```

```
end
```

3. Add a WCCP service group with service ID 90. This service group also specifies to cache sessions on ports 80 and 443 (for HTTP and HTTPS) and protocol number 6.

```
config system wccp
edit 90
set cache-id 10.51.101.10
set router-list 10.51.101.100
ports 80 443
set protocol 6
end
```

WCCP packet flow

The following packet flow sequence assumes you have configured a FortiGate unit to be a WCCP server and one or more FortiGate units to be WCCP clients.

1. A user's web browser sends a request for web content.
2. The FortiGate unit configured as a WCCP server includes a security policy that intercepts the request and forwards it to a WCCP client.

The security policy can apply UTM features to traffic accepted by the policy.

3. The WCCP client receives the WCCP session.
4. The client either returns requested content to the WCCP server if it is already cached, or connects to the destination web server, receives and caches the content and then returns it to the WCCP server.
5. The WCCP server returns the requested content to the user's web browser.
6. The WCCP router returns the request to the client web browser.

The client web browser is not aware that all this is taking place and does not have to be configured to use a web proxy.

Configuring the forward and return methods and adding authentication

The WCCP forwarding method determines how intercepted traffic is transmitted from the WCCP router to the WCCP cache engine. There are two different forwarding methods:

- GRE forwarding (the default) encapsulates the intercepted packet in an IP GRE header with a source IP address of the WCCP router and a destination IP address of the target WCCP cache engine. The result is a tunnel that allows the WCCP router to be multiple hops away from the WCCP cache server.
- L2 forwarding rewrites the destination MAC address of the intercepted packet to match the MAC address of the target WCCP cache engine. L2 forwarding requires that the WCCP router is Layer 2 adjacent to the WCCP client.

You can use the following command on a FortiGate unit configured as a WCCP router to change the forward and return methods to L2:

```
config system wccp
edit 1
set forward-method L2
set return-method L2
end
```

You can also set the forward and return methods to any in order to match the cache server configuration.

By default the WCCP communication between the router and cache servers is unencrypted. If you are concerned about attackers sniffing the information in the WCCP stream you can use the following command to enable hash-based authentication of the WCCP traffic. You must enable authentication on the router and the cache engines and all must have the same password.

```
config system wccp
  edit 1
    set authentication enable
    set password <password>
  end
```

WCCP messages

When the WCCP service is active on a web cache server it periodically sends a WCCP HERE I AM broadcast or unicast message to the FortiGate unit operating as a WCCP router. This message contains the following information:

- Web cache identity (the IP address of the web cache server).
- Service info (the service group to join).

If the information received in the previous message matches what is expected, the FortiGate unit replies with a WCCP I SEE YOU message that contains the following details:

- Router identity (the FortiGate unit's IP address).
- Sent to IP (the web cache IP addresses to which the packets are addressed)

When both ends receive these two messages the connection is established, the service group is formed and the designated web cache is elected.

Troubleshooting WCCP

Two types of debug commands are available for debugging or troubleshooting a WCCP connection between a FortiGate unit operating as a WCCP router and its WCCP cache engines.

Real time debugging

The following commands can capture live WCCP messages:

```
diag debug en
diag debug application wccpd <debug level>
```

Application debugging

The following commands display information about WCCP operations:

```
get test wccpd <integer>
diag test application wccpd <integer>
```

Where <integer> is a value between 1 and 6:

1. Display WCCP stats
2. Display WCCP config
3. Display WCCP cache servers
4. Display WCCP services
5. Display WCCP assignment
6. Display WCCP cache status

Enter the following command to view debugging output:

```
diag test application wccpd 3
```

Sample output from a successful WCCP connection:

```
service-0 in vdom-root: num=1, usable=1
cache server ID:
len=44, addr=172.16.78.8, weight=4135, status=0
rcv_id=6547, usable=1, fm=1, nq=0, dev=3(k3),
to=192.168.11.55
ch_no=0, num_router=1:
192.168.11.55
```

Sample output from the same command from an unsuccessful WCCP connection (because of a service group password mismatch):

```
service-0 in vdom-root: num=0, usable=0
diag debug application wccpd -1
Sample output:
wccp_on_recv()-98: vdom-root recv: num=160, dev=3(3),
172.16.78.8->192.168.11.55
wccp2_receive_pkt()-1124: len=160, type=10, ver=0200,
length=152
wccp2_receive_pkt()-1150: found component:t=0, len=20
wccp2_receive_pkt()-1150: found component:t=1, len=24
wccp2_receive_pkt()-1150: found component:t=3, len=44
wccp2_receive_pkt()-1150: found component:t=5, len=20
wccp2_receive_pkt()-1150: found component:t=8, len=24
wccp2_check_security_info()-326: MD5 check failed
```

Web proxy concepts

These are concepts that apply to both Transparent and Explicit Proxy.

Proxy policy

Information on Proxy policy options can be found at [Proxy option components on page 550](#)

Configuration information can be found at [Web proxy configuration on page 855](#)

Proxy authentication

Beginning in FortiOS 5.6, authentication is separated from authorization for user based policy. You can add authentication to proxy policies to control access to the policy and to identify users and apply different UTM features to different users. The described authentication methodology works with **Explicit Web Proxy** and **Transparent Proxy**.

Authentication of web proxy sessions uses HTTP basic and digest authentication as described in [RFC 2617 \(HTTP Authentication: Basic and Digest Access Authentication\)](#) and prompts the user for credentials from the browser allowing individual users to be identified by their web browser instead of IP address. HTTP authentication allows the FortiGate unit to distinguish between multiple users accessing services from a shared IP address.

The methodology of adding authentication has changed from FortiOS version 5.4 and previous version. Split-policy has been obsoleted and instead of identity-based-policy, authentication is managed by `authentication scheme`, `setting` and `rule` settings. These authentication settings are no longer configured with the individual policies. Authentication is set up in the contexts of:

```
config authentication scheme
config authentication setting
config authentication rule
```

The Authentication rule table defines how to identify user-ID. It uses the match factors:

- Protocol
- Source Address

For one address and protocol, there is only one authentication rule. It is possible to configure multiple authentication methods for on one address. The client browser will chose one authentication method from the authentication methods list, but you can not control which authentication method will be chosen by the browser.

Matching

If a rule is matched, the authentication methods defined in the rule will be used to authenticate a user. The procedure works as the following:

1. If it is IP-based, look up active user list to see a user existed from the source IP. If found, return the user ID.
2. If no method is set, an anonymous user is created to associate to the source-IP. Return the anonymous user. It is another way to bypass user authentication for some source IPs.
3. Use authentication methods to authenticate the user.
 - If no active method is defined, a failure will result to return an anonymous user.
 - Otherwise, a valid or guest user has to be identified to move on.

- Return the identified user ID.

Once a user is returned, the policy match resumes until a policy is matched or default policy will be used.

Processing policies for authentication

Authentication rules are checked once a User-ID is needed in order to resolve a match to a policy

Use the following scenario as an example of the process.

There are 3 policies:

- `policy1` does not have an associated user group
- `policy2` has an associated user group
- `policy3` does not have an associated user group

Step 1

If the traffic, based on protocol and source address matches `policy 1`, no user authentication is needed. The traffic is processed by `policy1`.

Step 2

If the traffic does not match `policy 1`, and any factor of `policy 2` is not matched, continue to next policy.

If all the factors except the user-group of `policy 2` are matched the authentication rule table is checked to get user-ID in the process in based on the procedure described earlier in Matching.

Step 3

When a user-ID is returned, whether it is a valid user or anonymous user, it is checked to see if the user is authorized by the user group associated with `policy2`. If yes, it is a match of `policy2`, and the traffic is processed by `policy2`. If not move on the next policy.

Step 4

For the purposes of the scenario, it will be assumed that the traffic either matches `policy3` or that `policy3` is the final policy that denies everything.

CLI syntax

Removals:

- "split-policy" from firewall explicit-proxy-policy.

The previous method to set up a split policy was:

```
config firewall explicit-proxy-policy
edit 1
    set proxy web
    set identity-based enable
    set groups <User group>
    config identity-based-policy
    edit 1
        set schedule "always"
        set utm-status enable
        set users "guest"
```



```

        set profile-protocol-options "default"
      next
    end
  next
end

```

- "auth relative" from firewall explicit-proxy-policy

The following attributes have been removed from firewall explicit-proxy-policy:

- identity-based
- ip-based
- active-auth-method
- sso-auth-method
- require-tfa

Moves:

users and groups from

```

firewall explicit-proxy-policy identity-based-policy
to

config firewall proxy-policy
  edit 1
    set groups <Group name>
    set users <User name>
  end

```

Additions:

authentication scheme

```

config authentication scheme
  edit <name>
    set method [ntlm|basic|digest|form|negotiate|fsso|rsso|none]
  end

```

- `ntlm` - NTLM authentication.
- `basic` - Basic HTTP authentication.
- `digest` - Digest HTTP authentication.
- `form` - Form-based HTTP authentication.
- `negotiate` - Negotiate authentication.
- `fsso` - FSSO authentication.
- `rsso` - RADIUS Single Sign-On authentication.
- `none` - No authentication.

authentication setting

```

config authentication setting
  set active-auth-scheme <string>
  set sso-auth-scheme <string>
  set captive-portal <string>
  set captive-portal-port <integer value from 1 to 65535>

```

- `active-auth-scheme` - Active authentication method.
- `sso-auth-scheme` - SSO authentication method.
- `captive-portal` - Captive portal host name.
- `captive-portal-port` - Captive portal port number.

authentication rule

```
config authentication rule
edit <name of rule>
    set status [enable|disable]
    set protocol [http|ftp|socks]
    set srcaddr <name of address object>
    set srcaddr6 <name of address object>
    set ip-based [enable|disable]
    set active-auth-method <string>
    set sso-auth-method <string>
    set web-auth-cookie [enable|disable]
    set transaction-based [enable|disable]
    set comments
```

- `status` - Enable/disable auth rule status.
- `protocol` - set protocols to be matched
- `srcaddr` / `srcaddr6` - Source address name. [`srcaddr` or `srcaddr6`(web proxy only) must be set].
- `ip-based` - Enable/disable IP-based authentication.
- `active-auth-method` - Active authentication method.
- `sso-auth-method` - SSO authentication method (require `ip-based` enabled)
- `web-auth-cookie` - Enable/disable Web authentication cookie.
- `transaction-based` - Enable/disable transaction based authentication.
- `comments` - Comment.

Configuring authentication in transparent proxy

You can enable transparent web-proxy feature to support authentication. Follow these steps

1. Configure a firewall policy
2. Enable a UTM profile in the firewall policy. Whenever there is a UTM item enabled, the feature enables the `profile-protocol-options`.
3. Go to the **Proxy Options** profile.
 - In the GUI this is **Security Profiles > Proxy Options**.
 - In the CLI it is `config firewall profile-protocol-options`.

Edit the profile used by the policy.

4. Enable HTTP in the profile.
In the GUI toggle on **HTTP** under **Protocol Port Mapping**

In the CLI, the command sequence is:

```
config firewall profile-protocol-options
edit <profile id>
    config http
        set status enable
    end
```

Fill out any other appropriate values.

5. Configure the proxy-policy, and set the value transparent-web for proxy option, others configuration are same as the explicit-web proxy

In the GUI, go to **Policy & Objects > Proxy Policy**. In the **Proxy Type** field choose **Transparent Web**.

In the CLI, the command sequence is:

```
config firewall proxy-policy
edit <profile id>
set proxy transparent-web
end
```

Fill out any other appropriate values.

6. Setup the authentication rule and scheme

With this configuration, if a HTTP request passes through FortiGate without explicit web proxy being applied, the traffic will be redirected to WAD daemon after it matches the proxy with HTTP-policy enabled, then WAD will do the proxy-policy matching, and all of the proxy authentication method can be used for the request.

Proxy addresses

Information on Proxy addresses can be found at [Proxy addresses on page 723](#)

Proxy address group

In the same way that IPv4 and IPv6 addresses can only be grouped together, Proxy addresses can only be grouped with other Proxy addresses. Unlike the other address groups, the Proxy address groups are further divided into source address groups and destination address groups. To see the configuration steps go to [Proxy address groups on page 725](#)

Web proxy firewall services and service groups

Configure web proxy services by selecting **Explicit Proxy** when configuring a service. Web proxy services can be selected in a explicit web proxy policy when adding one from the CLI. If you add a policy from the web-based manager the service is set to the **webproxy** service. The webproxy service should be used in most cases, it matches with any traffic with any port number. However, if you have special requirements, such as using a custom protocol type or a reduced port range or need to add an IP/FQDN to an proxy service you can create custom explicit web proxy services.

Web proxy services are similar to standard firewall services. You can configure web proxy services to define one or more protocols and port numbers that are associated with each web proxy service. Web proxy services can also be grouped into web proxy service groups.

One way in which web proxy services differ from firewall services is the protocol type you can select. The following protocol types are available:

- ALL
- CONNECT
- FTP
- HTTP

- SOCKS-TCP
- SOCKS-UDP

To add a web proxy service go to **Policy & Objects > Services** and select **Create New**. Set **Service Type** to **Explicit Proxy** and configure the service as required.

To add a web proxy service from the CLI enter:

```
config firewall service custom
edit my-socks-service
set explicit-proxy enable
set category Web Proxy
set protocol SOCKS-TCP
set tcp-portrange 3450-3490
end
```

To add a web proxy service group go to **Policy & Objects > Services** and select **Create New > Service Group**. Set **Type** to **Explicit Proxy** and add web proxy services to the group as required.

To add a web proxy service group from the CLI enter:

```
config firewall service group
edit web-group
set explicit-proxy enable
set member webproxy my-socks-service
end
```

Learn client IP

If there is another NATing device between the FortiGate and the Client (browser), this feature can be used to identify the real client in spite of the address translation. Knowing the actual client is imperative in cases where authorization is taking place.

The settings for the feature are in the CLI in the context of

```
config web-proxy global
```

Once here, enable the feature with the command:

```
set learn-client-ip enable
```

Once the feature is enabled, the other settings become available.

```
learn-client-ip-from-header
```

This command has the following options:

true-client-ip	Support HTTP header True-Client-IP.
x-real-ip	Support HTTP header X-Real-IP.
x-forwarded-for	Support HTTP header X-Forwarded-For.

```
learn-client-ip-srcaddr/learn-client-ip-srcaddr6
```

The options for this setting are selected from the list of IPv4 address or IPv6 address objects.

Example

Below is a config example where the real client ip address will be used to match policy or fsso authentication after the learn-client-ip feature enabled.

The value of `learn-client-ip-from-header` option can be set to `true-client-ip`, `x-real-ip` or `x-forwarded-for`, but in this case it has been set to `x-forward-for`.

```
config web-proxy global
    set proxy-fqdn "default.fqdn"
    set webproxy-profile "default"
    set learn-client-ip enable
    set learn-client-ip-from-header x-forwarded-for
    set learn-client-ip-srcaddr "all"
end

config firewall proxy-policy
    edit 1
        set proxy explicit-web
        set dstintf "mgmt1"
        set srcaddr "all"
        set dstaddr "all"
        set service "w"
        set action accept
        set schedule "always"
        set groups "fssol"
        set utm-status enable
        set av-profile "default"
        set dlp-sensor "default"
        set profile-protocol-options "default"
        set ssl-ssh-profile "deep-inspection"
    end

config authentication rule
    edit "rule1"
        set srcaddr "all"
        set sso-auth-method "scheme1"
    end

config authentication scheme
    edit "scheme1"
        set method fsso
    end
```

Web proxy configuration

General web proxy configuration steps

You can use the following general steps to configure the explicit web proxy.

To enable the explicit web proxy - web-based manager:

1. Go to **Network > Explicit Proxy** and enable **Explicit Web Proxy**. From here you can optionally change the HTTP port that the proxy listens on (the default is 8080) and optionally specify different ports for HTTPS, FTP, PAC, and other options.
2. Optionally enable **IPv6 Explicit Proxy** to turn on the explicit web proxy for IPv6 traffic.



If you enable both the IPv4 and the IPv6 explicit web proxy you can combine IPv4 and IPv6 addresses in a single explicit web proxy policy to allow both IPv4 and IPv6 traffic through the proxy.

3. Select **Apply**.
4. Go to **Network > Interfaces** and select one or more interfaces for which to enable the explicit web proxy. Edit the interface. Under the **Miscellaneous** heading select **Enable Explicit Web Proxy**.



Enabling the explicit web proxy on an interface connected to the Internet is a security risk because anyone on the Internet who finds the proxy could use it to hide their source address. If you enable the proxy on such an interface make sure authentication is required to use the proxy.

5. Go to **Policy & Objects > Addresses** and select **Create New** to add a firewall address that matches the source address of packets to be accepted by the explicit proxy.

Category	Address
Name	Internal_subnet
Type	IP Range
Subnet / IP Range	10.31.101.1 - 10.31.101.255
Interface	any*

*The **Interface** must be set to **Any**.

You can also set the **Type** to **URL Pattern (Explicit Proxy)** to add a destination URL that is only used by the explicit proxy. For example, to create an explicit policy that only allows access to Fortinet.com:

Category	Address
Name	Fortinet-web-sites
Type	URL Pattern (Explicit Proxy)
URL Pattern	fortinet.com
Interface	any

6. Go to **Policy & Objects > Proxy Policy** and select **Create New**. Configure the policy as required to accept the traffic that you want to be allowed to use the explicit web proxy.
7. Set the **Outgoing Interface** parameter by selecting the field with the "+" next to the field label. Selecting the field will slide out a window from the right where you can select from the available interfaces. You can select one or more specific interfaces. For more information on interfaces, check the Concepts section called [Interfaces and Zones](#).
8. The **Source** of the policy must match the client's source IP addresses. The interface of this firewall address must be set to **any**.
9. The **Destination** field should match the addresses of web sites that clients are connecting to. Usually the destination address would be **all** if proxying Internet web browsing. You could also specify a URL firewall address to limit the policy to allowing access to this URL.
10. Set the **Schedule** parameter by using the drop down menu to select a preconfigured schedule. The "+" icon next to the **Search** field is a shortcut for creating a new schedule object. For more information on addresses, check the Firewall Objects section called [Firewall schedules](#)
11. If **Default Firewall Policy Action** is set to **Deny** (under **Network > Explicit Proxy**), traffic sent to the explicit web proxy that is not accepted by a web-proxy policy is dropped. If **Default Firewall Policy Action** is set to **Allow** then all web-proxy sessions that don't match with a security policy are allowed.

For example, the following security policy allows users on an internal network to access fortinet.com websites through the wan1 interface of a FortiGate unit.

Explicit Proxy Type	Web
Source Address	Internal_subnet
Outgoing Interface	wan1
Destination Address	Fortinet-web-sites
Schedule	always
Action	ACCEPT



The explicit web-proxy accepts VIP addresses for destination addresses. If an external IP matches a VIP policy, the IP is changed to the mapped-IP of the VIP.

12. Set the **Disclaimer Options**

You can configure a disclaimer for each Authentication Rule by enabling one of the options here. The

choices are:

Disable	No disclaimer (default setting)
By Domain	The disclaimer will be displayed on different domains. The explicit web proxy will check the referring header to mitigate the javascript/css/images/video/etc page.
By Policy	The disclaimer will be displayed if the HTTP request matches a different explicit firewall policy.
By User	The disclaimer will be displayed when a new user logs on.

If you chose a disclaimer option other than **Disable**, you will have the option to enable **Customize Messages**. If enabled, select the **Edit Disclaimer Message** button to customize the message to your needs. This can be done as text or as HTML. The default HTML version is there if you just want to make minor changes.

13. Enable **Security Profiles** as required. Once the profile type is toggled to enabled, you can use the drop down menu to select a specific profile. The available profile types are:

- **AntiVirus**
- **WebFilter**
- **Application Control**
- **IPS**
- **DLP Sensor**
- **ICAP**
- **Web Application Firewall**

Just like with a regular policy, as soon as any of the **Security Profiles** is enabled, the following fields, with their own drop down menus for specific profiles will appear:

- **Proxy Options**
- **SSL/SSH Inspection**

14. Select **OK**.

To enable the explicit web proxy - CLI:

1. Enter the following command to turn on the IPv4 and IPv6 explicit web proxy for HTTP and HTTPS traffic.

```
config web-proxy explicit
  set status enable
  set ipv6-status enable
end
```

You can also enter the following command to enable the web proxy for FTP sessions in a web browser.

```
config web-proxy explicit
  set ftp-over-http enable
end
```

The default explicit web proxy configuration has `sec-default-action` set to `deny` and requires

you to add a security policy to allow access to the explicit web proxy.

2. Enter the following command to enable the explicit web proxy for the internal interface.

```
config system interface
edit internal
set explicit-web-proxy enable
end
end
```

3. Use the following command to add a firewall address that matches the source address of users who connect to the explicit web proxy.

```
config firewall address
edit Internal_subnet
set type iprange
set start-ip 10.31.101.1
set end-ip 10.31.101.255
end
```

The source address for a web-proxy security policy cannot be assigned to a FortiGate interface.

4. Optionally use the following command to add a destination URL that is only used by the explicit proxy. For example, to create an explicit policy that only allows access to Fortinet.com:

```
config firewall address
edit Fortinet-web-sites
set type url
set url fortinet.com
end
```

5. Use the following command to add an explicit web proxy policy that allows all users on the internal subnet to use the explicit web proxy for connections through the wan1 interface to the Internet.

```
config firewall proxy-policy
edit 0
set proxy explicit-web
set dstintf wan1
set scraddr Internal_subnet
set dstaddr all
set action accept
set service webproxy
set schedule always
end
```

6. Use the following command to add an explicit web proxy policy that allows authenticated users on the internal subnet to use the explicit web proxy for connections through the wan1 interface to the Internet.

```
config firewall proxy-policy
edit 0
set proxy explicit-web
set dstintf wan1
set scraddr Internal_subnet
set dstaddr Fortinet-web-sites
set action accept
set service webproxy
set schedule always
set groups <User group>
end
end
```

7. Use the following command to change global web proxy settings, for example to set the maximum request length for the explicit web proxy to 10:

```
config web-proxy global
  set max-request-length 10
end
```

8. Determine whether or not to use Botnet feature.

The option `scan-botnet-connections` uses the following syntax:

```
config firewall proxy-policy
  edit <policy id>
    set scan-botnet-connections [disable|block|monitor]
  end
```

Where:

- `disable` means do not scan connections to botnet servers
- `block` means block connection to botnet servers
- `monitor` means log connections to botnet servers

Logging options in web proxy profiles

There is an option on what action to take regarding the authenticated user's name in the header information for reading by upstream proxies and systems. This option can be used when a FortiGate is operating as an explicit proxy and authenticating users. The header is the `x-authenticated-user` and is used by the upstream proxy to ensure correct policy enforcement and to log the user's activity.

The `log-header-change` option enables the logging of any header changes in the web-proxy profile, including changes to authenticated users or groups.

Syntax

```
config web-proxy profile
  edit <profile ID#>
    set header-x-authenticated-user {pass|add|remove}
    set header-x-authenticated-groups {pass|add|remove}
    set log-header-change {enable|disable}
  end
```

Option	Description
header-x-authenticated-user	Action to take on the HTTP <code>x-authenticated-user</code> header in forwarded requests: <ul style="list-style-type: none">• <code>pass</code> - Forward the same HTTP header• <code>add</code> - Add the HTTP header• <code>remove</code> - Remove the HTTP header

Option	Description
header-x-authenticated-groups	<p>Action to take on the HTTP <code>x-authenticated-groups</code> header in forwarded requests:</p> <ul style="list-style-type: none"> • <code>pass</code> - Forward the same HTTP header • <code>add</code> - Add the HTTP header • <code>remove</code> - Remove the HTTP header
log-header-change	<code>enable</code> or <code>disable</code> the logging of HTTP header changes

Policy matching based on referrer headers and query strings

Web proxy policies support creating web proxy addresses to match referrer headers and query strings.

Matching referrer headers

For example, to create a web proxy address to match the referrer header to block access to the following YouTube URL `http://youtube.com/user/test321`. The http request will have the following format:

```
GET /user/test321 HTTP/1.1
Host: www.youtube.com
User-Agent: curl/7.52.1
Accept: */*
```

Create the following web proxy addresses to match this page:

```
config firewall proxy-address
edit youtube
set type host-regex
set host-regex ".*youtube.com"
next
edit test321
set host "youtube"
set path "/user/test321"
set referrer enable
end
```

Then create two proxy policies, one that allows access to all traffic and a second one that blocks access to the page that matches the referrer header:

```
config firewall proxy-policy
edit 1
set uuid 92273e4e-8c53-51e7-a7bd-f26e6e15fc98
set proxy explicit-web
set dstintf "wan2"
set srcaddr "all"
set dstaddr "all"
set service "webproxy-connect"
set action accept
set schedule "always"
set utm-status enable
set profile-protocol-options "test"
set ssl-ssh-profile "test"
next
edit 2
```

```

set uuid d35ad06a-8c53-51e7-8511-17200f682a4a
set proxy explicit-web
set dstintf "wan2"
set srcaddr "all"
set dstaddr "test321"
set service "webproxy"
set action accept
set schedule "always"
set utm-status enable
set av-profile "default"
set profile-protocol-options "test"
set ssl-ssh-profile "test"
end

```

Matching query strings

To match the video with URL `youtube.com/watch?v=XXXXXXXXXX`, (where `XXXXXXXXXX` is an example YouTube query string) you need to match an HTTP request with the following format:

```

GET /user/watch?v=GLCHldlwQsg HTTP/1.1
Host: www.youtube.com
User-Agent: curl/7.52.1
Accept: */*

```

Create the following web proxy addresses to match this video or query string:

```

config firewall proxy-address
edit "youtube"
set uuid 4ad63880-971e-51e7-7b2e-c69423ac6314
set type host-regex
set host-regex ".*youtube.com"
next
edit "query-string"
set uuid 7687a8c0-9727-51e7-5063-05edda03abbf
set host "youtube"
set path "/watch"
set query "v=XXXXXXXXXX"
end

```

Then create two proxy policies, one that allows access to all traffic and a second one that blocks access to the page that matches the query string

```

config firewall proxy-policy
edit 1
set uuid 92273e4e-8c53-51e7-a7bd-f26e6e15fc98
set proxy explicit-web
set dstintf "wan2"
set srcaddr "all"
set dstaddr "all"
set service "webproxy-connect"
set action accept
set schedule "always"
set utm-status enable
set profile-protocol-options "test"
set ssl-ssh-profile "test"
next
edit 2
set uuid d35ad06a-8c53-51e7-8511-17200f682a4a
set proxy explicit-web
set dstintf "wan2"

```

```
set srcaddr "all"
set dstaddr "query-string"
set service "webproxy"
set action accept
set schedule "always"
set utm-status enable
set av-profile "default"
set profile-protocol-options "test"
set ssl-ssh-profile "test"
next
end
```

Multiple web proxy PAC files in one VDOM

Proxy auto-config (PAC) files automatically choose the appropriate proxy server for browsers and other user agents. Not every user in an organization has the same proxy server requirements. Supporting multiple PAC files provides granular control. To manage multiple PAC files, you use PAC policies.

This capability is available only when the FortiGate is in **Proxy-based** inspection mode.

If there is no matching PAC policy (by name), in the PAC policies, the global PAC file is used by default.

To enable Proxy mode:

GUI

1. Go to **System > Settings**.
2. In **System Operation Settings**, set the **Inspection Mode** to **Proxy**.

CLI

```
config system settings
set inspection-mode proxy
end
```

To configure a PAC policy

```
config web-proxy explicit
set status enable
set pack-file-server-status enable
config pac-policy
edit <policy ID#>
set srcaddr <name of IPv4 address object>
set srcaddr6 <name of IPv6 address object>
set dstaddr <name of address object>
set pac-file-name <string>
set pac-file-data "<PAC-file>"
end
```

Option	Description
srcaddr or srcaddr6	This address must conform to the following criteria: <ul style="list-style-type: none">• a range, mask or wildcard mask type of address or address group• source type proxy-address or group It can be either IPv4 or IPv6.
dstaddr	This address must conform to the following criteria: <ul style="list-style-type: none">• a range, mask or wildcard type of address or address group• it must be resolved as the FortiGate address
pac-file-name	Name of the PAC file.
pac-file-data	Enter the contents of the PAC file enclosed in quotes. It is permissible to use the Return key when entering the contents. Place the closing quote at the end of the last line. If quotes are used within the content of the file, use the escape character \ before the quote. Example: \"

The `pac-file-server-status` setting must be set to `enable` in order for the `config pac-policy` command to work.

Explicit proxy concepts

The following is information that is specific to Explicit Proxy concepts. Any information that is common to Web Proxy in general is covered in the more inclusive section of [Web proxy concepts on page 848](#)

The FortiGate explicit web proxy

You can use the FortiGate explicit web proxy to enable explicit proxying of IPv4 and IPv6 HTTP, and HTTPS traffic on one or more FortiGate interfaces. The explicit web proxy also supports proxying FTP sessions from a web browser and proxy auto-config (PAC) to provide automatic proxy configurations for explicit web proxy users. From the CLI you can also configure the explicit web proxy to support SOCKS sessions from a web browser.

The explicit web and FTP proxies can be operating at the same time on the same or on different FortiGate interfaces.



If explicit web proxy options are not visible on the web-based manager, go to **System > Feature Visibility** and turn on **Explicit Proxy**.

In most cases you would configure the explicit web proxy for users on a network by enabling the explicit web proxy on the FortiGate interface connected to that network. Users on the network would configure their web browsers to use a proxy server for HTTP and HTTPS, FTP, or SOCKS and set the proxy server IP address to the IP address of the FortiGate interface connected to their network. Users could also enter the PAC URL into their web browser PAC configuration to automate their web proxy configuration using a PAC file stored on the FortiGate unit.



Enabling the explicit web proxy on an interface connected to the Internet is a security risk because anyone on the Internet who finds the proxy could use it to hide their source address.

If the FortiGate unit is operating in transparent mode, users would configure their browsers to use a proxy server with the FortiGate management IP address.

If the FortiGate unit is operating with multiple VDOMs the explicit web proxy is configured for each VDOM.

The web proxy receives web browser sessions to be proxied at FortiGate interfaces with the explicit web proxy enabled. The web proxy uses FortiGate routing to route sessions through the FortiGate unit to a destination interface. Before a session leaves the exiting interface, the explicit web proxy changes the source addresses of the session packets to the IP address of the exiting interface. When the FortiGate unit is operating in transparent mode the explicit web proxy changes the source addresses to the management IP address. You can configure the explicit web proxy to keep the original client IP address. See [The FortiGate explicit web proxy on page 864](#).

For more information about explicit web proxy sessions, see [The FortiGate explicit web proxy on page 864](#).

Example explicit web proxy topology



To allow all explicit web proxy traffic to pass through the FortiGate unit you can set the explicit web proxy default firewall policy action to accept. However, in most cases you would want to use security policies to control explicit web proxy traffic and apply security features such as access control/authentication, virus scanning, web filtering, application control, and traffic logging. You can do this by keeping the default explicit web proxy security policy action to deny and then adding web-proxy security policies.

You can also change the explicit web proxy default security policy action to accept and add explicit web proxy security policies. If you do this, sessions that match web-proxy security policies are processed according to the security policy settings. Connections to the explicit web proxy that do not match a web-proxy security policy are allowed with no restrictions or additional security processing. This configuration is not recommended and is not a best practice.

The explicit web-proxy can accept VIP addresses for destination address. If an external IP matches a VIP policy, the IP is changed to the mapped-IP of the VIP.

Web-proxy policies can selectively accept or deny traffic, apply authentication, enable traffic logging, and use security profiles to apply virus scanning, web filtering, IPS, application control, DLP, and SSL/SSH inspection to explicit web proxy traffic.

You cannot configure IPsec, SSL VPN, or Traffic shaping for explicit web proxy traffic. Web Proxy policies can only include firewall addresses not assigned to a FortiGate unit interface or with interface set to **Any**. (On the web-based manager you must set the interface to **Any**. In the CLI you must `unset the associated-interface`.)

Authentication of explicit web proxy sessions uses HTTP authentication and can be based on the user's source IP address or on cookies from the user's web browser. For more information, see [The FortiGate explicit web proxy on page 864](#).

To use the explicit web proxy, users must add the IP address of a FortiGate interface on which the explicit web proxy is enabled and the explicit web proxy port number (default 8080) to the proxy configuration settings of their web browsers.

On FortiGate units that support it, you can also enable web caching for explicit web proxy sessions.



For the time being, traffic shaping is not supported per policy for explicit proxy. For explicit proxy traffic, traffic shaping can be carried out per interface.

Other explicit web proxy options

You can change the following explicit web proxy options as required by your configuration.

HTTP port, HTTPS port, FTP port, PAC port

The TCP port that web browsers use to connect to the explicit proxy for HTTP, HTTPS, FTP and PAC services. The default port is 8080 for all services. By default HTTPS, FTP, and PAC use the same port as HTTP. You can change any of these ports as required. Users configuring their web browsers to use the explicit web proxy should add the same port numbers to their browser configurations.

Multi-port support for Explicit Proxy

Support exists for the use of multiple ports and port range in the explicit FTP or Web proxies. These changes have been added in both CLI and GUI.

CLI:

```
set http-incoming-port <port_low>[-<port_high>]
```

Where:

- `port_low` - the low value of the port
- `port_high` - the high value of the port

The `port_high` value can be omitted if `port_low` and `port_high` are the same.

Proxy FQDN

Enter the fully qualified domain name (FQDN) for the proxy server. This is the domain name to enter into browsers to access the proxy server.

Max HTTP request length

Enter the maximum length of an HTTP request in Kbytes. Larger requests will be rejected.

Max HTTP message length

Enter the maximum length of an HTTP message in Kbytes. Larger messages will be rejected.

Multiple incoming ports and port ranges

Web proxy can be configured to listen on multiple ports on the same IP as well as listen for HTTP and HTTPS on those same (or different) ports. This is done in the CLI.

Define the IP ranges using a hyphen (-). As shown below, `port_high` is not necessary to specify if `port_low` is equal to `port_high`.

CLI syntax

```
config web-proxy explicit
  set http-incoming-port <port_low> [-<port_high>]
end
```

Internet services

FortiOS can use the Internet Service Database (introduced in 5.4.1) as a web-proxy policy matching factor. This can only be done in the CLI.

CLI syntax:

```
config firewall proxy-policy
edit 0
set internet-service <application-id>
set internet-service-custom <application-name>
```

IP pools

IP Pools can be used with web proxy. When using this option of setting the IP pool name, the outgoing IP will be selected.

CLI syntax

```
config firewall proxy-policy
edit <example>
set poolname <name>
end
```

Proxy chaining (web proxy forwarding servers)

For the explicit web proxy you can configure web proxy forwarding servers to use proxy chaining to redirect web proxy sessions to other proxy servers. Proxy chaining can be used to forward web proxy sessions from the FortiGate unit to one or more other proxy servers on your network or on a remote network. You can use proxy chaining to integrate the FortiGate explicit web proxy with an web proxy solution that you already have in place.

A FortiGate unit can forward sessions to most web proxy servers including a remote FortiGate unit with the explicit web proxy enabled. No special configuration of the explicit web proxy on the remote FortiGate unit is required.

You can deploy the explicit web proxy with proxy chaining in an enterprise environment consisting of small satellite offices and a main office. If each office has a FortiGate unit, users at each of the satellite offices can use their local FortiGate unit as an explicit web proxy server. The satellite office FortiGate units can forward explicit web proxy sessions to an explicit web proxy server at the central office. From here the sessions can connect to web servers on the Internet.

FortiGate proxy chaining does not support authenticating with the remote forwarding server.

Adding a web proxy forwarding server

To add a forwarding server, select **Create New** in the **Web Proxy Forwarding Servers** section of the **Explicit Proxy** page by going to **Network > Explicit Proxy**.

Server Name	Enter the name of the forwarding server.
Proxy Address	Enter the IP address of the forwarding server.
Proxy Address Type	Select the type of IP address of the forwarding server. A forwarding server can have an FQDN or IP address.

Port	Enter the port number on which the proxy receives connections. Traffic leaving the FortiGate explicit web proxy for this server has its destination port number changed to this number.
Server Down action	<p>Select what action the explicit web proxy to take if the forwarding server is down.</p> <p>Block means if the remote server is down block traffic.</p> <p>Use Original Server means do not forward traffic to the forwarding sever but instead forward it from the FortiGate to its destination. In other words operate as if there is no forwarding server configured.</p>
Enable Health Monitor	Select to enable health check monitoring and enter the address of a remote site. See “Web proxy forwarding server monitoring and health checking” .
Health Check Monitor Site	

Use the following CLI command to add a web proxy forwarding server named `fwd-srv` at address `proxy.example.com` and port 8080.

```
config web-proxy forward-server
  edit fwd-srv
    set addr-type fqdn
    set fqdn proxy.example.com
    set port 8080
  end
```

Web proxy forwarding server monitoring and health checking

By default, a FortiGate unit monitors web proxy forwarding server by forwarding a connection to the remote server every 10 seconds. If the remote server does not respond it is assumed to be down. Checking continues and when the server does send a response the server is assumed to be back up. If you configure health checking, every 10 seconds the FortiGate unit attempts to get a response from a web server by connecting through the remote forwarding server.

You can configure health checking for each remote server and specify a different website to check for each one.

If the remote server is found to be down you can configure the FortiGate unit to block sessions until the server comes back up or to allow sessions to connect to their destination, bypassing the remote forwarding server. You cannot configure the FortiGate unit to fail over to another remote forwarding server.

Configure the server down action and enable health monitoring from the web-based manager by going to **Network > Explicit Proxy**, selecting a forwarding server, and changing the server down action and changing the health monitor settings.

Use the following CLI command to enable health checking for a web proxy forwarding server and set the server down option to bypass the forwarding server if it is down.

```
config web-proxy forward-server
  edit fwd-srv
    set healthcheck enable
    set monitor http://example.com
    set server-down-option pass
  end
```

Grouping forwarding servers and load balancing traffic to them

You can add multiple web proxy forwarding servers to a forwarding server group and then add the server group to an explicit web proxy policy instead of adding a single server. Forwarding server groups are created from the FortiGate CLI but can be added to policies from the web-based manager (or from the CLI).

When you create a forwarding server group you can select a load balancing method to control how sessions are load balanced to the forwarding servers in the server group. Two load balancing methods are available:

- **Weighted** load balancing sends more sessions to the servers with higher weights. You can configure the weight for each server when you add it to the group.
- **Least-session** load balancing sends new sessions to the forwarding server that is processing the fewest sessions.

When you create a forwarding server group you can also enable **affinity**. Enable affinity to have requests from the same client processed by the same server. This can reduce delays caused by using multiple servers for a single multi-step client operation. Affinity takes precedence over load balancing.

You can also configure the behavior of the group if all of the servers in the group are down. You can select to **block** traffic or you can select to have the traffic **pass** through the FortiGate explicit proxy directly to its destination instead of being sent to one of the forwarding servers.

Use the following command to add a forwarding server group that uses weighted load balancing to load balance traffic to three forwarding servers. Server weights are configured to send most traffic to server2. The group has affinity enabled and blocks traffic if all of the forward servers are down:

```
config web-proxy forward-server
  edit server_1
    set ip 172.20.120.12
    set port 8080
  next
  edit server_2
    set ip 172.20.120.13
    set port 8000
  next
  edit server_3
    set ip 172.20.120.14
    set port 8090
  next
end
config web-proxy forward-server-group
  edit New-fwd-group
    set affinity enable
    set ldb-method weight
    set group-down-option block
    config server-list
      edit server_1
        set weight 10
      next
      edit server_2
        set weight 40
      next
      edit server_3
        set weight 10
      next
    end
end
```

Adding proxy chaining to an explicit web proxy policy

You enable proxy chaining for web proxy sessions by adding a web proxy forwarding server or server group to an explicit web proxy policy. In a policy you can select one web proxy forwarding server or server group. All explicit web proxy traffic accepted by this security policy is forwarded to the specified web proxy forwarding server or server group.

To add an explicit web proxy forwarding server - web-based manager:

1. Go to **Policy & Objects > Proxy Policy** and select **Create New**.
2. Configure the policy:

Explicit Proxy Type	Web
Source Address	Internal_subnet
Outgoing Interface	wan1
Destination Address	all
Schedule	always
Action	ACCEPT
Web Proxy Forwarding Server	Select, fwd-srv

3. Select **OK** to save the security policy.

To add an explicit web proxy forwarding server - CLI:

1. Use the following command to add a security policy that allows all users on the 10.31.101.0 subnet to use the explicit web proxy for connections through the wan1 interface to the Internet. The policy forwards web proxy sessions to a remote forwarding server named `fwd-srv`

```
config firewall proxy-policy
edit 0
    set proxy explicit-web
    set dstintf wan1
    set scraddr Internal_subnet
    set dstaddr all
    set action accept
    set schedule always
    set webproxy-forward-server fwd-srv
end
```

Security profiles, threat weight, device identification, and the explicit web proxy

You can apply all security profiles to explicit web proxy sessions. This includes antivirus, web filtering, intrusion protection (IPS), application control, data leak prevention (DLP), and SSL/SSH inspection. Security profiles are applied by selecting them in an explicit web proxy policy or in authentication rules added to web proxy policies.

Traffic accepted by explicit web proxy policies contributes to threat weight data.

The explicit web proxy is not compatible with device identification.

Since the traffic accepted by the explicit web proxy is known to be either HTTP, HTTPS, or FTP over HTTP and since the ports are already known by the proxy, the explicit web proxy does not use all of the SSL/SSH inspection options. The explicit web proxy does support the following proxy options:

- Enable chunked bypass
- HTTP oversized file action and threshold

The explicit web proxy does not support the following proxy options:

- Client comforting
- Server comforting
- Monitor content information from dashboard. URLs visited by explicit web proxy users are not added to dashboard usage and log and archive statistics widgets.

For explicit web proxy sessions, the FortiGate unit applies antivirus scanning to HTTP POST requests and HTTP responses. The FortiGate unit starts virus scanning a file in an HTTP session when it receives a file in the body of an HTML request. The explicit web proxy can receive HTTP responses from either the originating web server or the FortiGate web cache module.

Explicit web proxy sessions and user limits

Web browsers and web servers open and close multiple sessions with the explicit web proxy. Some sessions open and close very quickly. HTTP 1.1 keepalive sessions are persistent and can remain open for long periods of time. Sessions can remain on the explicit web proxy session list after a user has stopped using the proxy (and has, for example, closed their browser). If an explicit web proxy session is idle for more than 3600 seconds it is torn down by the explicit web proxy. See [RFC 2616](#) for information about HTTP keepalive/persistent HTTP sessions.

This section describes proxy sessions and user limits for both the explicit web proxy and the explicit FTP proxy. Session and user limits for the two proxies are counted and calculated together. However, in most cases if both proxies are active there will be many more web proxy sessions than FTP proxy sessions.

The FortiGate unit adds two sessions to its session table for every explicit proxy session started by a web browser and every FTP session started by an FTP client. An entry is added to the session table for the session from the web browser or client to the explicit proxy. All of these sessions have the same destination port as the explicit web proxy port (usually 8080 for HTTP and 21 for FTP). An entry is also added to the session table for the session between the exiting FortiGate interface and the web or FTP server destination of the session. All of these sessions have a FortiGate interface IP address and the source address of the session and usually have a destination port of 80 for HTTP and 21 for FTP.

Proxy sessions that appear in FortiView do not include the Policy ID of the web-proxy or ftp-proxy security policy that accepted them. However, the explicit proxy sessions include a destination port that matches the explicit proxy port number (usually 8080 for the web proxy and 21 for the FTP proxy). The proxied sessions from the FortiGate unit have their source address set to the IP address of the FortiGate unit interface that the sessions use to connect to their destinations (for example, for connections to the Internet the source address would be the IP address of the FortiGate interface connected to the Internet).

FortiOS limits the number of explicit proxy users. This includes both explicit FTP proxy and explicit web proxy users. The number of users varies by FortiGate model from as low as 10 to up to 18000 for high end models. You cannot raise this limit.

If your FortiGate unit is configured for multiple VDOMs you can go to **System > Global Resources** to view the maximum number of **Concurrent explicit proxy users** and optionally reduce the limit. You can also use the following command:

```
config global
```

```
config system resource-limits
    set proxy 50
end
end
```

To limit the number of explicit proxy users for a VDOM, from the web-based manager enable multiple VDOMs and go to **System > VDOM** and edit a VDOM or use the following command to change the number of explicit web proxy users for VDOM_1:

```
config global
    config system vdom-property
        edit VDOM_1
            set proxy 25
        end
    end
end
```

You can use the `diagnose wad user list` command to view the number of explicit web proxy users. Users may be displayed with this command even if they are no longer actively using the proxy. All idle sessions time out after 3600 seconds.

You can use the command `diagnose wad user clear` to clear current explicit proxy users. You can also use the command `diagnose wad user clear <user-name>` to clear individual users. This means delete information about all users and force them re-authenticate.



Users that authenticate with explicit web-proxy or ftp-proxy security policies do not appear in the **Monitor > Firewall User Monitor** list and selecting **De-authenticate All Users** has no effect on explicit proxy users.

How the number of concurrent explicit proxy users is determined depends on their authentication method:

- For session-based authenticated users, each authenticated user is counted as a single user. Since multiple users can have the same user name, the proxy attempts to identify users according to their authentication membership (based upon whether they were authenticated using RADIUS, LDAP, FSAE, local database etc.). If a user of one session has the same name and membership as a user of another session, the explicit proxy assumes this is one user.
- For IP Based authentication, or no authentication, or if no web-proxy security policy has been added, the source IP address is used to determine a user. All sessions from a single source address are assumed to be from the same user.

The explicit proxy does not limit the number of active sessions for each user. As a result the actual explicit proxy session count is usually much higher than the number of explicit web proxy users. If an excessive number of explicit web proxy sessions is compromising system performance you can limit the amount of users if the FortiGate unit is operating with multiple VDOMs.

Explicit proxy configuration

The following is information that is specific to Explicit Proxy configuration. Any configuration information that is common to Web Proxy in general is covered in the more inclusive section of [Web proxy configuration on page 855](#).

Configuring an external IP address for the IPv4 explicit web proxy

You can use the following command to set an external IP address (or pool) that will be used by the explicit web proxy policy.

```
config web-proxy explicit
    set status enable
    set outgoing-ip <ip1> <ip2> ... <ipN>
end
```

Configuring an external IP address for the IPv6 explicit web proxy

You can use the following command to set an external IP address (or pool) that will be used by the explicit web proxy policy.

```
config web-proxy explicit
    set status enable
    set outgoing-ipv6 <ip1> <ip2> ... <ipN>
end
```

Restricting the IP address of the IPv4 explicit web proxy

You can use the following command to restrict access to the explicit web proxy using only one IP address. The IP address that you specify must be the IP address of an interface that the explicit web proxy is enabled on. You might want to use this option if the explicit FTP proxy is enabled on an interface with multiple IP addresses.

For example, to require uses to connect to the IP address 10.31.101.100 to connect to the explicit web proxy:

```
config web-proxy explicit
    set incoming-ip 10.31.101.100
end
```

Restricting the outgoing source IP address of the IPv4 explicit web proxy

You can use the following command to restrict the source address of outgoing web proxy packets to a single IP address. The IP address that you specify must be the IP address of an interface that the explicit web proxy is enabled on. You might want to use this option if the explicit web proxy is enabled on an interface with multiple IP addresses.

For example, to restrict the outgoing packet source address to 172.20.120.100:

```
config web-proxy explicit
    set outgoing-ip 172.20.120.100
end
```


Restricting the IP address of the explicit IPv6 web proxy

You can use the following command to restrict access to the IPv6 explicit web proxy to use only one IPv6 IP address. The IPv6 address that you specify must be the IPv6 address of an interface that the explicit web proxy is enabled on. You might want to use this option if the explicit web proxy is enabled on an interface with multiple IPv6 addresses.

For example, to require users to connect to the IPv6 address 2001:db8:0:2::30 to connect to the explicit IPv6 web proxy:

```
config web-proxy explicit
  set incoming-ipv6 2001:db8:0:2::30
end
```

Restricting the outgoing source IP address of the IPv6 explicit web proxy

You can use the following command to restrict the source address of outgoing web proxy packets to a single IPv6 address. The IP address that you specify must be the IPv6 address of an interface that the explicit web proxy is enabled on. You might want to use this option if the explicit web proxy is enabled on an interface with multiple IPv6 addresses.

For example, to restrict the outgoing packet source address to 2001:db8:0:2::50:

```
config web-proxy explicit
  set outgoing-ipv6 2001:db8:0:2::50
end
```

Explicit proxy firewall address types

Explicit proxy firewall address types improve granularity over header matching for explicit web proxy policies. You can enable this option using the **Show in Address List** button on the Address and Address Group New/Edit forms under **Policy & Objects > Addresses**.

The following address types are available:

- **URL Pattern** - destination address
- **Host Regex Match** - destination address
- **URL Category** - destination address (URL filtering)
- **HTTP Method** - source address
- **User Agent** - source address
- **HTTP Header** - source address
- **Advanced (Source)** - source address (combines User Agent, HTTP Method, and HTTP Header)
- **Advanced (Destination)** - destination address (combines Host Regex Match and URL Category)

Proxy auto-config (PAC) configuration

A proxy auto-config (PAC) file defines how web browsers can choose a proxy server for receiving HTTP content. PAC files include the FindProxyForURL(url, host) JavaScript function that returns a string with one or more access method specifications. These specifications cause the web browser to use a particular proxy server or to connect directly.

To configure PAC for explicit web proxy users, you can use the port that PAC traffic from client web browsers use to connect to the explicit web proxy. explicit web proxy users must configure their web browser's PAC proxy settings to use the PAC port.

PAC file content

You can edit the default PAC file from the web-based manager or use the following command to upload a custom PAC file:

```
config web-proxy explicit
    set pac-file-server-status enable
    set pac-file-data <pac_file_str>
end
```

Where <pac_file_str> is the contents of the PAC file. Enter the PAC file text in quotes. You can copy the contents of a PAC text file and paste the contents into the CLI using this option. Enter the command followed by two sets of quotes then place the cursor between the quotes and paste the file content.

The maximum PAC file size is 256 kbytes. If your FortiGate unit is operating with multiple VDOMs each VDOM has its own PAC file. The total amount of FortiGate memory available to store all of these PAC files 2 MBytes. If this limit is reached you will not be able to load any additional PAC files.

You can use any PAC file syntax that is supported by your users's browsers. The FortiGate unit does not parse the PAC file.

To use PAC, users must add an automatic proxy configuration URL (or PAC URL) to their web browser proxy configuration. The default FortiGate PAC file URL is:

```
http://<interface_ip>:<PAC_port_int>/<pac_file_str>
```

For example, if the interface with the explicit web proxy has IP address 172.20.120.122, the PAC port is the same as the default HTTP explicit web proxy port (8080) and the PAC file name is proxy.pac the PAC file URL would be:

```
http://172.20.120.122:8080/proxy.pac
```

From the CLI you can use the following command to display the PAC file URLs:

```
get web-proxy explicit
```

Unknown HTTP version

You can select the action to take when the proxy server must handle an unknown HTTP version request or message. Set unknown HTTP version to Reject or Best Effort. Best Effort attempts to handle the HTTP traffic as best as it can. Reject treats known HTTP traffic as malformed and drops it. The Reject option is more secure.

Authentication realm

You can enter an authentication realm to identify the explicit web proxy. The realm can be any text string of up to 63 characters. If the realm includes spaces enclose it in quotes. When a user authenticates with the explicit web proxy the HTTP authentication dialog includes the realm so you can use the realm to identify the explicitly web proxy for your users.

Implementing botnet features

The option `scan-botnet-connections` can be added to an explicit proxy policy.

CLI Syntax:

```
config firewall proxy-policy
```

```
edit <policy_id>
  set scan-botnet-connections [disable|block|monitor]
end
```

where:

- `disable` means do not scan connections to botnet servers.
- `block` means block connections to botnet servers.
- `monitor` means log connections to botnet servers.

Adding disclaimer messages to explicit proxy policies

This feature allows you to create user exceptions for specific URL categories (including warning messages) based on user groups. The **Disclaimer Options** are configured under **Policy & Objects > Proxy Policy**.

You can also configure a disclaimer for each Authentication Rule by setting **Action** to **Authenticate**.

Disclaimer explanations

- **Disable:** No disclaimer (default setting).
- **By Domain:** The disclaimer will be displayed on different domains. The explicit web proxy will check the referring header to mitigate the javascript/css/images/video/etc page.
- **By Policy:** The disclaimer will be displayed if the HTTP request matches a different explicit firewall policy.
- **By User:** The disclaimer will be displayed when a new user logs on.

Changing HTTP headers

You can create explicit web proxy profiles that can add, remove and change HTTP headers. The explicit web proxy profile can be added to a web explicit proxy policy and will be applied to all of the HTTP traffic accepted by that policy.

You can change the following HTTP headers:

- client-ip
- via header for forwarded requests
- via header for forwarded responses
- x-forwarded-for
- front-end-https

For each of these headers you can set the action to:

- Pass to forward the traffic without changing the header
- Add to add the header
- Remove to remove the header

You can also configure how the explicit web proxy handles custom headers. The proxy can add or remove custom headers from requests or responses. If you are adding a header you can specify the content to be included in the added header.

Create web proxy profiles from the CLI:

```
config web-proxy profile
edit <name>
  set header-client-ip {add | pass | remove}
```

```
set header-via-request {add | pass | remove}
set header-via-response {add | pass | remove}
set header-x-forwarded-for {add | pass | remove}
set header-front-end-https {add | pass | remove}
config headers
  edit <id>
    set action {add-to-request | add-to-response | remove-from-request |
              remove-from-response}
    set content <string>
    set name <name>
  end
end
```

Use the following command to add a web proxy profile to an explicit proxy policy:

```
config firewall proxy-policy
  edit <id>
    set webproxy-profile <name>
  end
```

Preventing the explicit web proxy from changing source addresses

By default in NAT/Route mode the explicit web proxy changes the source address of packets leaving the FortiGate to the IP address of the FortiGate interface that the packets are exiting from. In transparent mode the source address is changed to the management IP.

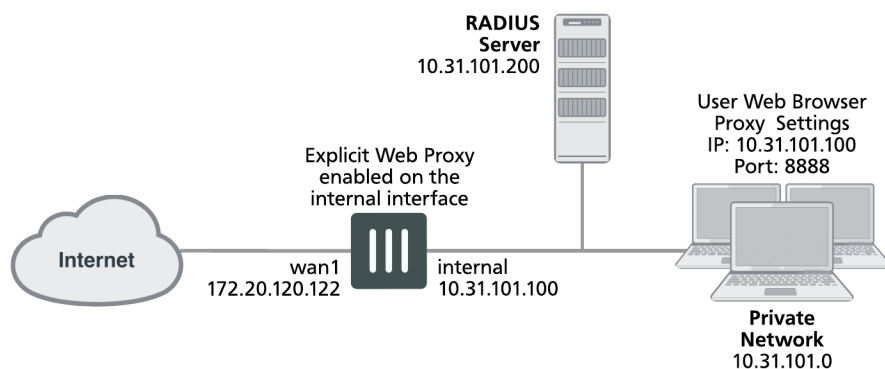
This configuration hides the IP addresses of clients and allows packets to return to the FortiGate unit interface without having to route packets from clients. You can use the following command to configure the explicit web proxy to keep the original client's source IP address:

```
config firewall proxy-policy
  edit 0
    set proxy explicit-web
    set transparent enable
  end
```

Example users on an internal network browsing the Internet through the explicit web proxy with web caching, RADIUS authentication, web filtering, and virus scanning

This example describes how to configure the explicit web proxy for the example network shown below. In this example, users on the internal network connect to the explicit web proxy through the Internal interface of the FortiGate unit. The explicit web proxy is configured to use port 8888 so users must configure their web browser proxy settings to use port 8888 and IP address 10.31.101.100.

Example explicit web proxy network topology



Explicit web proxy users must authenticate with a RADIUS server before getting access to the proxy. The explicit proxy policy that accepts explicit web proxy traffic applies per session authentication and includes a RADIUS server user group. The authentication rule also applies web filtering and virus scanning.

General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Enable the explicit web proxy for HTTP and HTTPS and change the HTTP and HTTPS ports to 8888.
2. Enable the explicit web proxy on the internal interface.
3. Add a RADIUS server and user group for the explicit web proxy.
4. Add an authentication explicit proxy policy. Enable web caching. Add an authentication rule and enable antivirus and web filtering.

Configuring the explicit web proxy - web-based manager

Use the following steps to configure the explicit web proxy.

To enable and configure the explicit web proxy

1. Go to **System > Feature Visibility** and turn on the **Explicit Proxy** feature.
2. Go to **Network > Explicit Proxy** and change the following settings:

Enable Explicit Web Proxy	Select HTTP/HTTPS .
Listen on Interfaces	No change. This field will eventually show that the explicit web proxy is enabled for the Internal interface.
HTTP Port	8888
HTTPS Port	0
Realm	You are authenticating with the explicit web proxy.
Default Firewall Policy Action	Deny

3. Select **Apply**.

To enable the explicit web proxy on the Internal interface

1. Go to **Network > Interfaces**.
2. Edit the internal interface.
3. Select **Enable Explicit Web Proxy**.
4. Select **OK**.

To add a RADIUS server and user group for the explicit web proxy

1. Go to **User & Device > RADIUS Servers** and select **Create New** to add a new RADIUS server:

Name	RADIUS_1
Primary Server Name/IP	10.31.101.200
Primary Server Secret	RADIUS_server_secret

2. Select **OK**.
3. Go to **User & Device > User Groups** and select **Create New** to add a new user group.

Name	Explicit_proxy_user_group
Type	Firewall
Remote Groups	RADIUS_1
Group Name	Any

4. Select **OK**.

To add an explicit proxy policy

1. Go to **Policy & Objects > Addresses** and select **Create New**.
2. Add a firewall address for the internal network:

Category	Address
Name	Internal_subnet
Type	Subnet / IP Range
Subnet / IP Range	10.31.101.0
Interface	Any

3. Go to **Policy & Objects > Proxy Policy** and select **Create New**.
4. Configure the explicit web proxy policy.

Explicit Proxy Type	Web
Source Address	Internal_subnet
Outgoing Interface	wan1
Destination Address	all
Action	AUTHENTICATE

- Under **Configure Authentication Rules** select **Create New** to add an authentication rule:

Groups	Explicit_policy
Source User(s)	Leave blank
Schedule	always

- Turn on **Antivirus** and **Web Filter** and select the **default** profiles for both.
- Select the **default** proxy options profile.
- Select **OK**.
- Make sure **Enable IP Based Authentication** is not selected.
- Turn on **Web Cache**.
- Select **OK**.

Configuring the explicit web proxy - CLI

Use the following steps to configure the example explicit web proxy configuration from the CLI.

To enable the explicit web proxy on the Internal interface

- Enter the following command to enable the explicit web proxy on the internal interface.

```
config system interface
  edit internal
    set explicit-web-proxy enable
  end
```

To enable and configure the explicit web proxy

- Enter the following command to enable the explicit web proxy and set the TCP port that proxy accepts HTTP and HTTPS connections on to 8888.

```
config web-proxy explicit
  set status enable
  set http-incoming-port 8888
  set https-incoming-port 8888
  set realm "You are authenticating with the explicit web proxy"
  set sec-default-action deny
end
```

To add a RADIUS server and user group for the explicit web proxy

- Enter the following command to add a RADIUS server:

```
config user radius
```

```
edit RADIUS_1
  set server 10.31.101.200
  set secret RADIUS_server_secret
end
```

2. Enter the following command to add a user group for the RADIUS server.

```
config user group
  edit Explicit_proxy_user_group
    set group-type firewall
    set member RADIUS_1
  end
```

To add a security policy for the explicit web proxy

1. Enter the following command to add a firewall address for the internal subnet:

```
config firewall address
  edit Internal_subnet
    set type iprange
    set start-ip 10.31.101.1
    set end-ip 10.31.101.255
  end
```

2. Enter the following command to add the explicit web proxy security policy:

```
config firewall proxy-policy
  edit 0
    set proxy explicit-web
    set dstintf wan1
    set srcaddr Internal_subnet
    set dstaddr all
    set action accept
    set service webproxy
    set webcache enable
    set identity-based enable
    set ipbased disable
    set active-auth-method basic
    set groups <User group>
  end
```

Testing and troubleshooting the configuration

You can use the following steps to verify that the explicit web proxy configuration is working as expected:

To test the explicit web proxy configuration

1. Configure a web browser on the internal subnet to use a web proxy server at IP address 10.31.101.100 and port 8888.
2. Browse to an Internet web page.
The web browser should pop up an authentication window that includes the phrase that you added to the Realm option.
3. Enter the username and password for an account on the RADIUS server.
If the account is valid you should be allowed to browse web pages on the Internet.
4. Close the browser and clear its cache and cookies.
5. Restart the browser and connect to the Internet.
You could also start a second web browser on the same PC. Or you could start a new instance of the same

browser as long as the browser asks for a user name and password again.

You should have to authenticate again because identity-based policies are set to session-based authentication.

6. If this basic functionality does not work, check your FortiGate and web browser configuration settings.
7. Browse to a URL on the URL filter list and confirm that the web page is blocked.
8. Browse to <http://eicar.org> and attempt to download an anti-malware test file.

The antivirus configuration should block the file.

Sessions for web-proxy security policies do not appear on the Top Sessions dashboard widget and the count column for security policies does not display a count for explicit web proxy security policies.

9. You can use the following command to display explicit web proxy sessions

```
get test wad 60
IP based users:

Session based users:
  user:0x9c20778, username:User1, vf_id:0, ref_cnt:9

Total allocated user:1

Total user count:3, shared user quota:50, shared user count:3
```

This command output shows one explicit proxy user with user name `User1` authenticated using session-based authentication.

Kerberos and NTLM authentication

FortiOS recognizes the client's authentication method from the token and selects the correct authentication scheme to authenticate successfully.

CLI syntax

```
config firewall proxy-policy
  edit 0
    set active-auth-method [ntlm|basic|digest|negotiate|none]
  end
```

Kerberos authentication for explicit proxy users

Kerberos authentication is a method for authenticating both explicit web proxy and transparent web proxy users. It has several advantages over NTLM challenge response:

- Does not require FSSO/AD agents to be deployed across domains.
- Requires fewer round-trips than NTLM SSO, making it less latency sensitive.
- Is (probably) more scalable than challenge response.
- Uses existing Windows domain components rather than added components.
- NTLM may still be used as a fallback for non-Kerberos clients.

Enhancements to Kerberos explicit and transparent web proxy

FortiOS 5.6.x authentication is managed by schemes and rules based on protocol and source address. As such, configurable authentication settings have been introduced to enhance authentication.

CLI commands (`config authentication rule, scheme, and setting`) allow explicit proxy rules and schemes to be created to separate user authentication (e.g. authentication rules and schemes used to match conditions in order to identify users) from user authorization (proxy-based policies with users and/or user groups).

CLI syntax - config authentication rule

```
config authentication rule
edit <name>
    set name <name>
    set status {enable|disable}
    set protocol {http|ftp|socks}
    config srcaddr <addr-name or addrgrp-name>
        edit <name>
            set name <ipv4-policy-name>
        next
    end
    config srcaddr6 <addr-name or addrgrp-name>
        edit <name>
            set name <ipv6-policy-name>
        next
    end
    set ip-based {enable|disable}
    set active-auth-method <scheme-name>
    set sso-auth-method <scheme-name>
    set transaction-based {enable|disable} - basic scheme + session-based
    set web-auth-cookie {enable|disable}
    set comments <comments>
next
end
```

Note: As shown above, HTTP, FTP, and SOCKSV5 authentication protocols are supported for explicit proxy.

Authentication rules are used to receive user-identity, based on the values set for protocol and source address. Having said this, if a rule fails to match based on source address, there will be no other attempt to match the rule, however the next policy will be attempted. This occurs only when:

- there is an authentication rule, but no authentication method has been set (under `config authentication scheme`; see below), so user identity cannot be found.
- the user is successfully matched in the rule, but fails to match the current policy.

Once a rule is positively matched through protocol and/or source address, it must also match the authentication method specified (`active-auth-method` and `sso-auth-method`). These methods point to schemes, as defined under `config authentication scheme`.

CLI syntax - config authentication scheme

```
config authentication scheme
edit <name>
    set name <name>
    set method {basic|digest|ntlm|form|negotiate|fsso|rsso}
    set negotiate-ntlm {enable|disable}
    set require-tfa {enable|disable}
    set fsso-guest {enable|disable}
    config user-database
        edit <name>
            set name {local|<ldap-server>|<radius-server>|<fsso-name>|<rsso-name>|<tacacs+-name>}
        next
    end
end
```

```

        next
    end
    next
end

```

Combining authentication rules and schemes, granular control can be exerted over users and IPs, creating an efficient process for users to successfully match a criteria before matching the policy.

Additional options can be set under `config authentication setting`.

CLI syntax - config authentication setting

```

config authentication setting
    set sso-scheme <scheme-name>
    set active-scheme <scheme-name>
    set captive-portal <host-name>
    set captive-portal-port <tcp-port>
end

```

Integration of transparent and explicit proxy HTTP policy checking

A CLI command, under `config firewall profile-protocol-options`, allows HTTP policy checking to be enable or disabled. When enabled, transparent traffic can be matched in a firewall policy and policy user authentication can occur. In addition, separate SSL inspection policies can be created:

```

config firewall profile-protocol-options
    edit <name>
        set http-policy {enable|disable}
    end

```

Internet Service Database in Explicit/Implicit proxy policies

CLI commands, under `config firewall proxy-policy`, implement the Internet Service Database (ISDB) as the webproxy matching factor, and override IP pool is also support:

```

config firewall proxy-policy
    edit <name>
        set proxy {explicit-web|transparent-web|ftp|wanopt}
        set dstintf <dst-name>
        set poolname <ip-pool-name>
    end

```

Multiple port/port range support for explicit web and explicit FTP proxy

Multiple port numbers and/or ranges can be set for explicit proxy, specifically for HTTP/HTTPS and FTP. Go to **Network > Explicit Proxy** and configure settings under **Explicit Web Proxy** and **Explicit FTP Proxy**, or under `config web-proxy explicit` in the CLI Console.

1. General configuration

1.1 Kerberos environment - Windows server setup

1. Build a Windows 2008 Platform server.
2. Enable domain configuration in windows server (dcpromo).
3. Set the domain name TEST.COM (realm name).

1.2 Create users

- *testuser* is a normal user (could be any existing domain user account).
- *testfgt* is the service name. In this case it should be the FQDN for the explicit proxy Interface, For example the hostname in the client browser proxy config.
- Recommendation: create username all in lowercase (even if against corporate standards).
 - The account only requires “domain users” membership
 - Password set to never expire
 - Set a very strong password

1.3 Add FortiGate to DNS



Add the FortiGate FQDN in to the Windows DNS domain, as well as in-addr.arpa

For Lab/Testing add the FortiGate Domain name and IP mapping in the hosts file (windows/system32/drivers/etc/hosts). e.g., `TESTFGT.TEST.COM 10.10.1.10`

1.4 Generate the Kerberos keytab

Use the *ktpass* command (found on Windows Servers and many domain workstations) to generate the Kerberos keytab.

Example:

```
ktpass -princ HTTP/<domain name of test fgt>@realm -mapuser testfgt -pass <password> -crypto all -ptype KRB5_NT_PRINCIPAL -out fgt.keytab
```



In the case where the FortiGate is handling multiple keytabs in Kerberos authentication, use different passwords when generating each keytab.



The ktpass on older Windows servers (i.e. 2003) may not support the “all” crypto option.

Example:

```
ktpass -princ HTTP/testfgt.test.com@TEST.COM -mapuser testfgt -pass 12345678 -crypto all -ptype KRB5_NT_PRINCIPAL -out fgt.keytab
```



The realm name is always presented in uppercase, and prefixed with the “@” character.

1.5 Encode base64

Use the *base64* command (available in most Linux distros) command to encode the *fgt.keytab* file. Any LF (Line Feed) need to be deleted from the file.

Example:

```
base64 fgt.keytab > fgt.txt
```



Use Notepad++ or some native Linux text editor. Windows Notepad and Wordpad are likely to introduce errors.

2. FortiGate configuration

2.1 Create LDAP server instance

```
config user ldap
  edit "ldap" <<< Required for authorization
    set server "10.10.1.1" <<< LDAP server IP, normally it should be same as KDC server
    set cnid "cn"
    set dn "dc=test,dc=com"
    set type regular
    set username "CN=admin,CN=Users,DC=test,DC=com" <<< Your domain may require STARTTLS
    set password <FOOS>
  next
end
```

2.2 Define Kerberos as an authentication service

```
config user krb-keytab
  edit "http_service"
    set principal "HTTP/testfgt.test.com@TEST.COM" <<< Same as the principal name in 1.4
    set ldap-server "ldap" <<< the defined ldap server for authorization
    set keytab
      "BQIAAABNAAIACKJFUKJFUi5DT00ABEhUVFAAGlRPTllfRkdUXzEwMERfQS5CRVJCRVIuQ09NAAAAQA
      AAAAKABcAEJQl0MHgovwplu7XzfENJzw=" <<< base64 encoding keytab data, created in step 1.5
  next
end
```

2.3 Create user group(s)

```
config user group <<< the group is used for kerberos authentication
  edit "testgrp"
    set member "ldap"
    config match
      edit 1
        set server-name "ldap" <<< Same as ldap-server option in krb-keytab
        set group-name "CN=Domain Users,CN=Users,DC=TEST,DC=com"
      next
    end
  next
end
```

2.4 Create firewall policy

```
config firewall proxy-policy
  edit 1
    set uuid 5e5dd6c4-952c-51e5-b363-120ad77c1414
    set proxy explicit-web
    set dstintf "port1"
```

```

        set srcaddr "all"
        set dstaddr "all"
        set service "webproxy"
        set action accept
        set schedule "always"
        set groups "CN=USERS LAB.PS FSSO"
    next
end

```

2.5 Diagnostics

Once the keytab is imported, check that it has been properly decoded. The filename generated will be relatively random, but should be clearly visible.

```

Artoo-Deetoo (root) # fnsysctl ls -la /tmp/kt
drwxr--r-- 2 0 0 Fri Dec 2 10:06:43 2016 60 .
drwxrwxrwt 22 0 0 Tue Dec 6 14:28:29 2016 3280 ..
-rw-r--r-- 1 0 0 Fri Dec 2 10:06:43 2016 392 1.0.89.keytab

```



If there is no file present, then the file hasn't decoded. Check the file for line feeds and try again.

3. Client side walkthrough

3.1 Check Kerberos is working

Log on to the domain by using *testuser*, created in 1.2. Use the *klist* command to list ticket information. In the below example, the client has received *krbtgt*, *CIFS*, and *LDAP* tickets. As there has been no interaction with the FortiGate, there are no references to it.

```

C:\Users\glenk>klist
Cached Tickets: (5)

C:\Users\glenk>klist
Cached Tickets: (5)
#0> Client: glenk @ home.local

    Server: krbtgt/HOME.LOCAL @ HOME.LOCAL
    KerbTicket Encryption Type: AES-256-CTS-HMAC-SHA1-96
    Ticket Flags 0x60a00000 -> forwardable forwarded renewable pre_authent
    Start Time: 12/6/2016 14:58:06 (local)
    End Time: 12/7/2016 0:58:04 (local)
    Renew Time: 12/13/2016 14:58:04 (local)
    Session Key Type: AES-256-CTS-HMAC-SHA1-96

#1> Client: glenk @ home.local

    Server: krbtgt/HOME.LOCAL @ HOME.LOCAL
    KerbTicket Encryption Type: AES-256-CTS-HMAC-SHA1-96
    Ticket Flags 0x40e00000 -> forwardable renewable initial pre_authent
    Start Time: 12/6/2016 14:58:04 (local)
    End Time: 12/7/2016 0:58:04 (local)
    Renew Time: 12/13/2016 14:58:04 (local)
    Session Key Type: AES-256-CTS-HMAC-SHA1-96

#2> Client: glenk @ home.local

    Server: cifs/EthicsGradient.home.local @ HOME.LOCAL

```

```
KerbTicket Encryption Type: AES-256-CTS-HMAC-SHA1-96
Ticket Flags 0x40a40000 -> forwardable renewable pre_authent ok_as_delegate
Start Time: 12/6/2016 14:58:06 (local)
End Time: 12/7/2016 0:58:04 (local)
Renew Time: 12/13/2016 14:58:04 (local)
Session Key Type: AES-256-CTS-HMAC-SHA1-96
```

```
#3> Client: glenk @ home.local
```

```
Server: ldap/EthicsGradient.home.local @ HOME.LOCAL
KerbTicket Encryption Type: AES-256-CTS-HMAC-SHA1-96
Ticket Flags 0x40a40000 -> forwardable renewable pre_authent ok_as_delegate
Start Time: 12/6/2016 14:58:06 (local)
End Time: 12/7/2016 0:58:04 (local)
Renew Time: 12/13/2016 14:58:04 (local)
Session Key Type: AES-256-CTS-HMAC-SHA1-96
```

```
#4> Client: glenk @ home.local
```

```
Server: LDAP/EthicsGradient.home.local/home.local @ HOME.LOCAL
KerbTicket Encryption Type: AES-256-CTS-HMAC-SHA1-96
Ticket Flags 0x40a40000 -> forwardable renewable pre_authent ok_as_delegate
Start Time: 12/6/2016 14:58:06 (local)
End Time: 12/7/2016 0:58:04 (local)
Renew Time: 12/13/2016 14:58:04 (local)
Session Key Type: AES-256-CTS-HMAC-SHA1-96
```

3.2 Configure client

Set up web-proxy in browser through the FortiGate. This can be achieved via a PAC file or direct browser configuration.



Some Firefox documentation indicates that it is necessary to make manual advanced configuration changes to allow Kerberos authentication work. However, builds 48 (and possibly much earlier) require no additional configuration beyond setting of the proxy server.

3.3 Open a connection to the Internet

1. The client accesses the explicit proxy, but a *HTTP 407 Proxy Authentication Required* is returned.
2. As “Negotiate” is set, the client has knowledge of the KRBtgt, it requests a ticket from the KDC with a *krb-tgs-req* message. This includes the REALM (HOME.LOCAL) in the *req-body* section, and the provided instances SNAME and service (in this case, HTTP/artoo-deetoo.home.local).
3. The KDC responds with a next KRB-TGS-REP.

This ticket is then available on the client.

In the example below, the ticket-granted-service has issued Ticket #2.

```
#2> Client: glenk @ home.local
Server: HTTP/artoo-deetoo.home.local @ HOME.LOCAL
KerbTicket Encryption Type: RSADSI RC4-HMAC (NT)
Ticket Flags 0x40a00000 -> forwardable renewable pre_authent
Start Time: 12/6/2016 14:59:45 (local)
End Time: 12/7/2016 0:58:04 (local)
Renew Time: 12/13/2016 14:58:04 (local)
Session Key Type: RSADSI RC4-HMAC (NT)
```

4. The conversation between the client and the proxy continues, as the client responds with the Kerberos ticket in the response.

The whole process takes less than a second to complete. The user should be visible as a FSSO logon in the Web UI.

Transparent web-proxy Kerberos authentication

Transparent web-proxy allows the FortiGate to process level 7 policy matching, even when the explicit web-proxy is not enabled on the client's browser. The transparent web-proxy policy is set in proxy-policy too. The policy matching rule is the same as the explicit web-proxy.

In the firewall policy level, transparent web-proxy is regarded as a special UTM. The HTTP/HTTPS traffic matches the firewall policy first, then traffic is redirected to the web-proxy daemon. If the transparent web-proxy feature is disabled, http-policy options in profile-protocol-options is used to enable transparent web-proxy feature.

IP-based

Kerberos authentication requires the captive portal to be an FQDN address that is resolved to a local IP address. However, it becomes more complicated to setup an FQDN address in a local user deployment. Therefore you can set the `captive-portal-type` to either use an FQDN or IP address.

1. Captive portal and the captive portal port must be configured in transparent web-proxy for support of Kerberos authentication:

```
config authentication setting
    set captive-portal-type {fqdn | ip}
    set captive-portal <fqdn-name> / <ip>
    set captive-portal-port "9998"
end
```

2. Authentication rule, scheme, and krb-keytab need to be configured for Kerberos authentication (note the `active-auth-method` scheme referenced in the rule):

```
config authentication scheme
    edit <kerberos-scheme>
        set method negotiate
        set negotiate-ntlm <enable>
        set fsso-guest <disable>
    next
end

config authentication rule
    edit <name>
        set status <enable>
        set protocol <http>
        set srcaddr "all"
        set ip-based <enable>
        set active-auth-method <kerberos-scheme>
    next
end

config user krb-keytab
    edit <name>
        set principal "HTTP/TESTFGT.TEST.COM@TEST.COM"
        set ldap-server "ldap"
```



```

        set keytab <base64-encoding-keytab-data>
    next
end

```

3. Configure LDAP and user group used for authorization:

```

config user ldap
    edit "ldap"
        set server "10.10.1.1"
        set cnid
        srt dn
        set type <regular>
        set username "CN=admin,CN=Users,DC=test,DC=com"
        set password ENC
            aW51IAHkPMf4D+ZCKpGMU3x8Fpq0G+7uIbAvpblbXFA5vLfGb4/oRBx+B6R/v+CMCetP84e+Gdz5zEcM
            yOd3cj0BoIhFrpYJfXhRs4lSEOHezeVXfxwTSf5VJG+F11G/G5RpaY+AE8bortC8MBe7P2/uGQocFHu4
            Ilulp5I60Jvyk6Ei3hDZMjTd8iPp5IkRJZVVjQ==
    next
end

config user group
    edit "testgrp"
        set member "ldap"
        config match
            edit "1"
                set server-name "ldap"
                set group-name "CN=Domain Users,CN=Users,DC=TEST,DC=com"
            next
        end
    next
end

```

4. Create proxy-policy, with groups as the authorizing policy-matching element:

```

config firewall proxy-policy
    edit 1
        set uuid 1bbb891a-9cd2-51e7-42ff-d1fa13cac3da
        set proxy explicit-web
        set dstintf "any"
        set srcaddr "all"
        set dstaddr "all"
        set service "webproxy"
        set action accept
        set schedule "always"
        set groups testgrp
    next
end

```

5. UTM must be enabled in the firewall policy to support the transparent web-proxy:

```

config firewall policy
    edit "1"
        set name "policy1"
        set uuid 8a6ceeac-b016-51e6-2b5c-165070d5bf50
        set srcintf "mgmt1"
        set dstintf "mgmt1"
        set srcaddr "all"
        set dstaddr "all"
        set action <accept>
    next
end

```

```
        set schedule "always"
        set service "ALL"
        set utm-status <enable>
        set profile-protocol-options "transparent-web-proxy"
        set ssl-ssh-profile "deep-inspection"
        set nat <enable>
    next
end

config firewall profile-protocol-options
    edit "transparent-web-proxy"
        config http
            set ports "80 8080"
            unset options
            set http-policy enable
            unset post-lang
        end
        ...
    next
end
```

Session-based with web-auth cookie

The web-auth-cookie feature is necessary for session-based authentication under transparent web-proxy.

The configuration is the same as for IP-based authentication, except `ip-based` is disabled in the authentication rule:

```
config authentication rule
    edit "kerberos-rules"
        set status <enable>
        set protocol <http>
        set srcaddr "all"
        set ip-based <disable>
        set active-auth-method <kerberos-scheme>
    next

config authentication setting
    set captive-portal <fqdn-name>
    set captive-portal-port "9998"
end
```

Transparent proxy concepts

In addition to the Explicit Web Proxy, FortiOS supports a Transparent web proxy. While it does not have as many features as Explicit Web Proxy, the transparent proxy has the advantage that nothing needs to be done on the user's system to forward supported web traffic over to the proxy. There is no need to reconfigure the browser or publish a PAC file. Everything is transparent to the end user, hence the name. This makes it easier to incorporate new users into a proxy deployment.

You can use the transparent proxy to apply web authentication to HTTP traffic accepted by a firewall policy. In previous versions of FortiOS, web authentication required using the explicit proxy.

Normal FortiOS authentication is IP address based. Users are authenticated according to their IP address and access is allowed or denied based on this IP address. On networks where authentication based on IP address will not work you can use the Transparent Web proxy to apply web authentication that is based on the user's browser and not on their IP address. This authentication method allows you to identify individual users even if multiple users on your network are connecting to the FortiGate from the same IP address.

More about the transparent proxy

The following changes are incorporated into Transparent proxy, some of which affect Explicit Web Proxy as well.

Flat policies

The split policy feature has been removed. This will make the explicit policy more like the firewall policy.

Authentication

The authentication design is intended to separate authentication from authorization. Authentication has been moved into a new table in the FortiOS. This leaves the authorization as the domain of the explicit proxy policy.

Previously, if authentication was to be used:

1. The policy would be classified as an identity based policy
2. The policy would be split to add the authentication parameters
3. The authentication method would be selected
4. The user/group would be configured

Now:

The user/group is configured in the proxy policy

1. A new authentication rule is added
2. This option refers to the authentication scheme
3. The authentication scheme has the details of the authentication method

The new authentication work flow for transparent proxy:

Toggle the transparent-http-policy match:

```
config firewall profile-protocol-options
edit <profile ID>
config http
```

```
set http-policy <enable|disable>
```

If disabled, everything works like before. If enabled, the authentication is triggered differently.

- http-policy work flow:
- For transparent traffic, if there is a regular firewall policy match, when the Layer 7 check option is enabled, traffic will be redirected to WAD for further processing.
- For redirected traffic, layer 7 policy (HTTP policy) will be used to determine how to do security checks.
- If the last matching factor is down to user ID, then it will trigger a new module to handle the L7 policy user authentication.
- Then propagate learned user information back to the system so that it can be used to match traffic for L4 policy.

New proxy type

There is a new subcategory of proxy in the proxy policy called **Transparent Web**. The old **Web Proxy** is now referred to as **Explicit Web Proxy**.

- This is set in the firewall policy
- It is available when the HTTP policy is enabled in the profile-protocol options for the firewall policy
- This proxy type supports OSI layer 7 address matching.
- This proxy type should include a source address as a parameter
- Limitations:
 - It can be used for HTTPS traffic, if deep scanning is not used
 - It only supports SNI address matching, i.e. domain names
 - It does not support header types of address matching
 - It only supports SSO authentication methods, no active authentication methods.

IP pools support

Proxies are now supported on outgoing IP pools.

SOCKSv5

SOCKSv5 authentication is now supported for explicit proxies.

To configure:

```
config authentication rule
  edit <name of rule>
    set protocol socks
  end
```

Forwarding

Proxies support URL redirect/forwarding. This allows a non-proxy forwarding server to be assigned a rule that will redirect web traffic from one URL to another, such as redirecting traffic destined for youtube.com to restrict.youtube.com.

- A new option called "Redirect URL" has been added to the policy
- Traffic forwarding by VIP is supported

Support for explicit proxy address objects & groups into IPv4 firewall policies

This would allow the selection of web filter policy, SSL inspection policy, and proxy policy based on source IP + destination (address|explicit proxy object|category|group of any of those). This enables things like “do full SSL interception on www.google.com, but not the rest of the Search Engines category”.

Support application service in the proxy based on HTTP requests.

The application service can be configured using the following CLI commands:

```
config firewall service custom
  edit <name of service>
    set explicit-proxy enable
    set app-service-type <disable|app-id|app-category>
    set app-category <application category ID, integer>
    set application <application ID, integer>
  end
```

Transparent proxy configuration

To implement the Transparent proxy, go to **System > Settings** and scroll down to **Operations Settings** and set the inspection mode to **Proxy**.

Operations Settings

Inspection Mode

Flow-based

Proxy

Virtual Domains ☐

Then go to **System > Feature Visibility** and enable **Explicit Proxy**.

Security Features

Feature Set: Custom

☒ Anti-Spam Filter +

☒ AntiVirus +

☒ Application Control +

☒ DLP +

☒ DNS Filter +

☒ Endpoint Control +

☒ Explicit Proxy +

☒ Intrusion Prevention +

☒ Web Application Firewall +

☒ Web Filter +

Then go to **Security Profiles > Proxy Options**, edit a proxy options profile and under **Web Options** enable **HTTP Policy Redirect**.

Web Options

Chunked Bypass ☐

Add Fortinet Bar ☐

HTTP Policy Redirect ☒

Then go to **Policy & Objects > IPv4 Policy** and create or edit a policy that accepts traffic that you want to apply web authentication to. This can be a general policy that accepts many different types of traffic as long as it also accepts the web traffic that you want to apply web authentication to.

Select a **Security Profile** and select the **Proxy Options** profile that you enabled **HTTP Policy Redirect** for.

Name	General Internet Access Policy		
Incoming Interface	port2	▼	
Outgoing Interface	port1	▼	
Source	all	✕	
Destination	all	✕	
Schedule	always	▼	
Service	ALL	✕	
Action	<input checked="" type="checkbox"/> ACCEPT	<input type="checkbox"/> DENY	<input type="checkbox"/> LEARN

Firewall / Network Options

NAT ☒






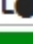
IP Pool Configuration ☒ Use Outgoing Interface Address ☐ Use Dynamic IP Pool

Security Profiles


AntiVirus	<input checked="" type="checkbox"/>	default	✎
Web Filter	<input type="checkbox"/>		
DNS Filter	<input type="checkbox"/>		
Application Control	<input type="checkbox"/>		
IPS	<input type="checkbox"/>		
Proxy Options		default	✎
SSL/SSH Inspection		certificate-inspection	✎

Then go to **Policy & Objects > Proxy Policy** create a Transparent Proxy policy to accept the traffic that you want to apply web authentication to. Set the **Proxy Type** to **Transparent Web**. The incoming interface, outgoing interface, destination address, and schedule should either match or be a subset of the same options defined in the IPv4 policy. Addresses added to the Source must match or be a subset of the source addresses added to the IPv4 policy. You can also add the users to be authenticated by the transparent policy to the source field.

Select other transparent policy options as required.

Proxy Type 	Explicit Web Transparent Web FTP
Incoming Interface	 port2 ▼
Outgoing Interface	 port1 ▼
Source	 web_users ✕
Destination Address	 all ✕
Schedule	 always ▼
Action	<input checked="" type="checkbox"/> ACCEPT <input type="checkbox"/> DENY

Disclaimer Options

Display Disclaimer	Disable By Domain By Policy By User
Customize Messages <input checked="" type="checkbox"/>	 Edit Disclaimer Message

Security Profiles

AntiVirus	<input type="checkbox"/>
Web Filter	<input type="checkbox"/>
Application Control	<input type="checkbox"/>
IPS	<input type="checkbox"/>
Web Proxy Forwarding Server	<input type="checkbox"/>

Logging Options

Log Allowed Traffic <input checked="" type="checkbox"/>	Security Events All Sessions
Comments	<input type="text" value="Write a comment..."/> 0/1023

CLI changes due to addition of transparent proxy

The adding of Transparent Proxy to the existing proxy types has required some changes, removals, moves and additions to the CLI.

Changes:

Previous	New
<code>config firewall explicit-proxy-policy</code>	<code>config firewall proxy-policy</code>
<code>config firewall explicit-proxy-address</code>	<code>config firewall proxy-address</code>
<code>config firewall explicit-proxy-addrgrp</code>	<code>config firewall proxy-addrgrp</code>
<pre>config firewall explicit-proxy-policy edit <policy ID> set proxy web end</pre>	<pre>config firewall proxy-policy edit <policy ID> set proxy explicit-web end</pre>

FTP proxy concepts

The FortiGate explicit FTP proxy

You can use the FortiGate explicit FTP proxy to enable explicit FTP proxying on one or more FortiGate interfaces. The explicit web and FTP proxies can be operating at the same time on the same or on different FortiGate interfaces.



Explicit FTP proxies are configured for each VDOM when multiple VDOMs are enabled.

In most cases you would configure the explicit FTP proxy for users on a network by enabling the explicit FTP proxy on the FortiGate interface connected to that network. Users on the network would connect to and authenticate with the explicit FTP proxy before connecting to an FTP server. In this case the IP address of the explicit FTP proxy is the IP address of the FortiGate interface on which the explicit FTP proxy is enabled.




Enabling the explicit FTP proxy on an interface connected to the Internet is a security risk because anyone on the Internet who finds the proxy could use it to hide their source address.

If the FortiGate unit is operating in transparent mode, users would configure their browsers to use a proxy server with the FortiGate unit management IP address.

The FTP proxy receives FTP sessions to be proxied at FortiGate interfaces with the explicit FTP proxy enabled. The FTP proxy uses FortiGate routing to route sessions through the FortiGate unit to a destination interface. Before a session leaves the exiting interface, the explicit FTP proxy changes the source addresses of the session packets to the IP address of the exiting interface. When the FortiGate unit is operating in transparent mode the explicit web proxy changes the source addresses to the management IP address.

Example explicit FTP proxy topology



To allow anyone to anonymously log into explicit FTP proxy and connect to any FTP server you can set the explicit FTP proxy default firewall proxy action to accept. When you do this, users can log into the explicit FTP proxy with any username and password.

In most cases you would want to use explicit proxy policies to control explicit FTP proxy traffic and apply security features, access control/authentication, and logging. You can do this by keeping the default explicit FTP proxy firewall policy action to deny and then adding explicit FTP proxy policies. In most cases you would also want users to authenticate with the explicit FTP proxy. By default an anonymous FTP login is required. Usually you would add authentication to explicit FTP proxy policies. Users can then authenticate with the explicit FTP proxy

according to users or user groups added to the policies. User groups added to explicit FTP proxy policies can use any authentication method supported by FortiOS including the local user database and RADIUS and other remote servers.

If you leave the default firewall policy action set to deny and add explicit FTP proxy policies, all connections to the explicit FTP proxy must match an or else they will be dropped. Sessions that are accepted are processed according to the ftp-proxy security policy settings.

You can also change the explicit FTP proxy default firewall policy action to accept and add explicit FTP proxy policies. If you do this, sessions that match explicit FTP proxy policies are processed according to the policy settings. Connections to the explicit FTP proxy that do not match an explicit FTP proxy policy are allowed and the users can authenticate with the proxy anonymously.

There are some limitations to the security features that can be applied to explicit FTP proxy sessions. See [The FortiGate explicit FTP proxy on page 899](#).

You cannot configure IPsec, SSL VPN, or Traffic shaping for explicit FTP proxy traffic. Explicit FTP proxy policies can only include firewall addresses not assigned to a FortiGate unit interface or with interface set to **any**. (On the web-based manager you must set the interface to **Any**. In the CLI you must `unset the associated-interface`.)

How to use the explicit FTP proxy to connect to an FTP server

To connect to an FTP server using the explicit FTP proxy, users must run an FTP client and connect to the IP address of a FortiGate interface on which the explicit FTP proxy is enabled. This connection attempt must use the configured explicit FTP proxy port number (default 21).

The explicit FTP proxy is not compatible with using a web browser as an FTP client. To use web browsers as FTP clients configure the explicit web proxy to accept FTP sessions.

The following steps occur when a user starts an FTP client to connect to an FTP server using the explicit FTP proxy. Any RFC-compliant FTP client can be used. This example describes using a command-line FTP client. Some FTP clients may require a custom FTP proxy connection script.

1. The user enters a command on the FTP client to connect to the explicit FTP proxy.

For example, if the IP address of the FortiGate interface on which the explicit FTP proxy is enabled is 10.31.101.100, enter:

```
ftp 10.31.101.100
```

2. The explicit FTP proxy responds with a welcome message and requests the user's FTP proxy user name and password and a username and address of the FTP server to connect to:

```
Connected to 10.31.101.100.
220 Welcome to FortiGate FTP proxy
Name (10.31.101.100:user):
```

You can change the message by editing the FTP Explicit Banner Message replacement message.

3. At the prompt the user enters their FTP proxy username and password and a username and address for the FTP server. The FTP server address can be a domain name or numeric IP address. This information is entered using the following syntax:

```
<proxy-user>:<proxy-password>:<server-user>@<server-address>
```

For example, if the proxy username and password are `p-name` and `p-pass` and a valid username for the FTP server is `s-name` and the server's IP address is `ftp.example.com` the syntax would be:

```
p-name:p-pass:s-name@ftp.example.com
```

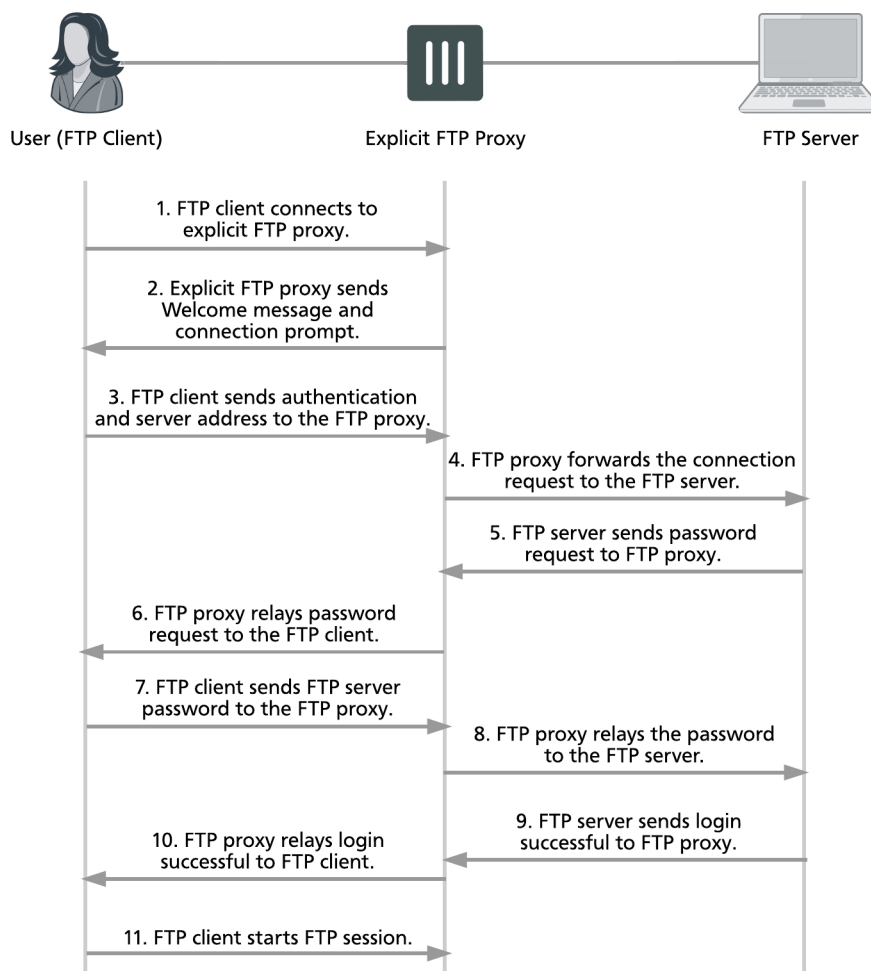


If the FTP proxy accepts anonymous logins `p-name` and `p-pass` can be any characters.

4. The FTP proxy forwards the connection request, including the user name, to the FTP server.
5. If the user name is valid for the FTP server it responds with a password request prompt.
6. The FTP proxy relays the password request to the FTP client.
7. The user enters the FTP server password and the client sends the password to the FTP proxy.
8. The FTP proxy relays the password to the FTP server.
9. The FTP server sends a login successful message to the FTP proxy.
10. The FTP proxy relays the login successful message to the FTP client.
11. The FTP client starts the FTP session.

All commands entered by the client are relayed by the proxy to the server. Replies from the server are relayed back to the FTP client.

Explicit FTP proxy session



From a simple command line FTP client connecting to an the previous sequence could appear as follows:

```
ftp 10.31.101.100 21
Connected to 10.31.101.100.
220 Welcome to FortiGate FTP proxy
Name (10.31.101.100:user): p-name:p-pass:s-name@ftp.example.com
331 Please specify the password.
Password: s-pass
230 Login successful.
Remote system type is UNIX
Using binary mode to transfer files.
ftp>
```

Security profiles, threat weight, device identification, and the explicit FTP proxy

You can apply antivirus, data leak prevention (DLP), and SSL/SSH inspection to explicit FTP proxy sessions. Security profiles are applied by selecting them in an explicit FTP proxy policy or an authentication rule in an FTP proxy security policy.

Traffic accepted by explicit FTP proxy policies contributes to threat weight data.

The explicit FTP proxy is not compatible with device identification.

Explicit FTP proxy options and SSL/SSH inspection

Since the traffic accepted by the explicit FTP proxy is known to be FTP and since the ports are already known by the proxy, the explicit FTP proxy does not use the FTP port proxy options settings.

When adding UTM features to an FTP proxy security policy, you must select a proxy options profile. In most cases you can select the default proxy options profile. You could also create a custom proxy options profile.

The explicit FTP proxy supports the following proxy options:

- Block Oversized File and oversized file limit

The explicit FTP proxy does not support the following protocol options:

- Client comforting

Explicit FTP proxy sessions and antivirus

For explicit FTP proxy sessions, the FortiGate unit applies antivirus scanning to FTP file GET and PUT requests. The FortiGate unit starts virus scanning a file in an FTP session when it receives a file in the body of an FTP request.

Flow-based virus scanning is not available for explicit FTP proxy sessions. Even if the FortiGate unit is configured to use flow-based antivirus, explicit FTP proxy sessions use the regular virus database.

Explicit FTP proxy sessions and user limits

FTP clients do not open large numbers of sessions with the explicit FTP proxy. Most sessions stay open for a short while depending on how long a user is connected to an FTP server and how large the file uploads or downloads are. So unless you have large numbers of FTP users, the explicit FTP proxy should not be adding large numbers of sessions to the session table.

Explicit FTP proxy sessions and user limits are combined with explicit web proxy session and user limits. For information about explicit proxy session and user limits, see [Explicit proxy sessions and user limits on page 1](#).

FTP proxy configuration

General explicit FTP proxy configuration steps

You can use the following general steps to configure the explicit FTP proxy.

To enable the explicit FTP proxy - web-based manager:

1. Go to **Network > Explicit Proxy > Explicit FTP Proxy Options**. Select **Enable Explicit FTP Proxy** to turn on the explicit FTP proxy.
2. Select **Apply**.

The **Default Firewall Policy Action** is set to **Deny** and requires you to add a explicit FTP proxy policy to allow access to the explicit FTP proxy. This configuration is recommended and is a best practice because you can use policies to control access to the explicit FTP proxy and also apply security features and authentication.

3. Go to **Network > Interfaces** and select one or more interfaces for which to enable the explicit web proxy. Edit the interface and select **Enable Explicit FTP Proxy**.



Enabling the explicit FTP proxy on an interface connected to the Internet is a security risk because anyone on the Internet who finds the proxy could use it to hide their source address. If you enable the proxy on such an interface make sure authentication is required to use the proxy.

4. Go to **Policy & Objects > Proxy Policy** and select **Create New** and set the **Explicit Proxy Type** to **FTP**.

You can add multiple explicit FTP proxy policies.

5. Configure the policy as required to accept the traffic that you want to be processed by the explicit FTP proxy.

The source address of the policy should match client source IP addresses. The firewall address selected as the source address cannot be assigned to a FortiGate interface. The Interface field of the firewall address must be blank or it must be set to **Any**.

The destination address of the policy should match the IP addresses of FTP servers that clients are connecting to. The destination address could be **all** to allow connections to any FTP server.

If **Default Firewall Policy Action** is set to Deny, traffic sent to the explicit FTP proxy that is not accepted by an explicit FTP proxy policy is dropped. If **Default Firewall Policy Action** is set to Allow then all FTP proxy sessions that don't match a policy are allowed.

For example the following explicit FTP proxy policy allows users on an internal network to access FTP servers on the Internet through the wan1 interface of a FortiGate unit.

Explicit Proxy Type	FTP
Source Address	Internal_subnet

Outgoing Interface	wan1
Destination Address	all
Schedule	always
Action	ACCEPT

The following explicit FTP proxy policy requires users on an internal network to authenticate with the FortiGate unit before accessing FTP servers on the Internet through the wan1 interface.

Explicit Proxy Type	FTP
Source Address	Internal_subnet
Outgoing Interface	wan1
Destination Address	all
Action	AUTHENTICATE

6. Select **Create New** to add an **Authentication Rule** and configure the rule as follows:

Groups	Proxy-Group
Source Users	(optional)
Schedule	always

7. Add security profiles as required and select **OK**.
8. You can add multiple authentication rules to apply different authentication for different user groups and users and also apply different security profiles and logging settings for different users.
9. Select **OK**.

To enable the explicit FTP proxy - CLI:

1. Enter the following command to turn on the explicit FTP proxy. This command also changes the explicit FTP proxy port to 2121.

```
config ftp-proxy explicit
  set status enable
  set incoming-port 2121
end
```

The default explicit FTP proxy configuration has `sec-default-action` set to `deny` and requires you to add a security policy to allow access to the explicit FTP proxy.

2. Enter the following command to enable the explicit FTP proxy for the internal interface.

```
config system interface
  edit internal
    set explicit-ftp-proxy enable
  end
end
```

3. Use the following command to add a firewall address that matches the source address of users who connect to the explicit FTP proxy.

```
config firewall address
  edit Internal_subnet
    set type iprange
    set start-ip 10.31.101.1
    set end-ip 10.31.101.255
  end
```

The source address for a ftp-proxy security policy cannot be assigned to a FortiGate unit interface.

4. Use the following command to add an explicit FTP proxy policy that allows all users on the internal subnet to use the explicit FTP proxy for connections through the wan1 interface to the Internet.

```
config firewall proxy-policy
  edit 0
    set proxy ftp
    set dstintf wan1
    set scraddr Internal_subnet
    set dstaddr all
    set action accept
    set schedule always
  end
```

5. Use the following command to add an explicit FTP proxy policy that allows authenticated users on the internal subnet to use the explicit FTP proxy for connections through the wan1 interface to the Internet.

```
config firewall proxy-policy
  edit 0
    set proxy ftp
    set dstintf wan1
    set scraddr Internal_subnet
    set dstaddr Fortinet-web-sites
    set action accept
    set schedule always
    set groups <User group>
  end
end
```

Restricting the IP address of the explicit FTP proxy

You can use the following command to restrict access to the explicit FTP proxy using only one IP address. The IP address that you specify must be the IP address of an interface that the explicit FTP proxy is enabled on. You might want to use this option if the explicit FTP proxy is enabled on an interface with multiple IP addresses.

For example, to require users to connect to the IP address 10.31.101.100 to connect to the explicit FTP proxy:

```
config ftp-proxy explicit
  set incoming-ip 10.31.101.100
end
```

Restricting the outgoing source IP address of the explicit FTP proxy

You can use the following command to restrict the source address of outgoing FTP proxy packets to a single IP address. The IP address that you specify must be the IP address of an interface that the explicit FTP proxy is enabled on. You might want to use this option if the explicit FTP proxy is enabled on an interface with multiple IP addresses.


For example, to restrict the outgoing packet source address to 172.20.120.100:

```
config ftp-proxy explicit
  set outgoing-ip 172.20.120.100
end
```

Example users on an internal network connecting to FTP servers on the Internet through the explicit FTP with RADIUS authentication and virus scanning

This example describes how to configure the explicit FTP proxy for the example network shown below. In this example, users on the internal network connect to the explicit FTP proxy through the Internal interface with IP address 10.31.101.100. The explicit web proxy is configured to use port 2121 so to connect to an FTP server on the Internet users must first connect to the explicit FTP proxy using IP address 10.31.101.100 and port 2121.

Example explicit FTP proxy network topology



In this example, explicit FTP proxy users must authenticate with a RADIUS server before getting access to the proxy. To apply authentication, the security policy that accepts explicit FTP proxy traffic includes an identity based policy that applies per session authentication to explicit FTP proxy users and includes a user group with the RADIUS server in it. The identity based policy also applies UTM virus scanning and DLP.

General configuration steps

This section breaks down the configuration for this example into smaller procedures. For best results, follow the procedures in the order given:

1. Enable the explicit FTP proxy and change the FTP port to 2121.
2. Enable the explicit FTP proxy on the internal interface.
3. Add a RADIUS server and user group for the explicit FTP proxy.
4. Add a user identity security policy for the explicit FTP proxy.
5. Enable antivirus and DLP features for the identity-based policy.

Configuring the explicit FTP proxy - web-based manager

Use the following steps to configure the explicit FTP proxy from FortiGate web-based manager.

To enable and configure the explicit FTP proxy

1. Go to **Network > Explicit Proxy > Explicit FTP Proxy Options** and change the following settings:

Enable Explicit FTP Proxy	Select.
Listen on Interface	No change. This field will eventually show that the explicit web proxy is enabled for the Internal interface.
FTP Port	2121
Default Firewall Policy Action	Deny

2. Select **Apply**.

To enable the explicit FTP proxy on the Internal interface

1. Go to **Network > Interfaces**, edit the Internal interface and select **Enable Explicit FTP Proxy**.

To add a RADIUS server and user group for the explicit FTP proxy

1. Go to **User & Device > RADIUS Servers**.
2. Select **Create New** to add a new RADIUS server:

Name	RADIUS_1
Primary Server Name/IP	10.31.101.200
Primary Server Secret	RADIUS_server_secret

3. Go to **User > User > User Groups** and select **Create New**.

Name	Explicit_proxy_user_group
Type	Firewall
Remote groups	RADIUS_1
Group Name	ANY

4. Select **OK**.

To add a security policy for the explicit FTP proxy

1. Go to **Policy & Objects > Addresses** and select **Create New**.
2. Add a firewall address for the internal network:

Address Name	Internal_subnet
Type	Subnet
Subnet / IP Range	10.31.101.0
Interface	Any

3. Go to **Policy & Objects > Proxy Policy** and select **Create New**.
4. Configure the explicit FTP proxy security policy.

Explicit Proxy Type	FTP
Source Address	Internal_subnet
Outgoing Interface	wan1
Destination Address	all
Action	AUTHENTICATE

5. Under **Configure Authentication Rules** select **Create New** to add an authentication rule:

Groups	Explicit_policy
Users	Leave blank
Schedule	always

6. Turn on **Antivirus** and **Web Filter** and select the **default** profiles for both.
7. Select the **default** proxy options profile.
8. Select **OK**.
9. Make sure **Enable IP Based Authentication** is not selected and **DefaultAuthentication Method** is set to **Basic**.
10. Select **OK**.

Configuring the explicit FTP proxy - CLI

Use the following steps to configure the example explicit web proxy configuration from the CLI.

To enable and configure the explicit FTP proxy

1. Enter the following command to enable the explicit FTP proxy and set the TCP port that proxy accepts FTP connections on to 2121.

```
config ftp-proxy explicit
  set status enable
  set incoming-port 2121
  set sec-default-action deny
end
```

To enable the explicit FTP proxy on the Internal interface

1. Enter the following command to enable the explicit FTP proxy on the internal interface.

```
config system interface
  edit internal
    set explicit-ftp-proxy enable
  end
```

To add a RADIUS server and user group for the explicit FTP proxy

1. Enter the following command to add a RADIUS server:

```
config user radius
edit RADIUS_1
set server 10.31.101.200
set secret RADIUS_server_secret
end
```

2. Enter the following command to add a user group for the RADIUS server.

```
config user group
edit Explicit_proxy_user_group
set group-type firewall
set member RADIUS_1
end
```

To add a security policy for the explicit FTP proxy

1. Enter the following command to add a firewall address for the internal subnet:

```
config firewall address
edit Internal_subnet
set type iprange
set start-ip 10.31.101.1
set end-ip 10.31.101.255
end
```

2. Enter the following command to add the explicit FTP proxy security policy:

```
config firewall proxy-policy
edit 0
set proxy ftp
set dstintf wan1
set srcaddr Internal_subnet
set dstaddr all
set action accept
set identity-based enable
set ipbased disable
set active-auth-method basic
set groups <User group>
end
```

Testing and troubleshooting the configuration

You can use the following steps to verify that the explicit FTP proxy configuration is working as expected. These steps use a command line FTP client.

To test the explicit web proxy configuration

1. From a system on the internal network start an FTP client and enter the following command to connect to the FTP proxy:

```
ftp 10.31.101.100
```

The explicit FTP proxy should respond with a message similar to the following:

```
Connected to 10.31.101.100.
220 Welcome to Floodgate FTP proxy
Name (10.31.101.100:user):
```

2. At the prompt enter a valid username and password for the RADIUS server followed by a user name for an FTP server on the Internet and the address of the FTP server. For example, if a valid username and password on the

RADIUS server is ex_name and ex_pass and you attempt to connect to an FTP server at ftp.example.com with user name s_name, enter the following at the prompt:

```
Name (10.31.101.100:user):ex_name:ex_pass:s_name@ftp.example.com
```

3. You should be prompted for the password for the account on the FTP server.
4. Enter the password and you should be able to connect to the FTP server.
5. Attempt to explore the FTP server file system and download or upload files.
6. To test UTM functionality, attempt to upload or download an ECAR test file. Or upload or download a text file containing text that would be matched by the DLP sensor.

For eicar test files, go to <http://eicar.org>.

Diagnose commands for WAN optimization

The following get and diagnose commands are available for troubleshooting WAN optimization, web cache, explicit proxy and WCCP.

get test {wad | wccpd} <test_level>

Display usage information about WAN optimization, explicit proxy, web cache, and WCCP applications. Use <test_level> to display different information.

```
get test wad <test_level>
get test wccpd <test_level>
```

Variable	Description
wad	Display information about WAN optimization, web caching, the explicit web proxy, and the explicit FTP proxy.
wccpd	Display information about the WCCP application.

Examples

Enter the following command to display WAN optimization tunnel protocol statistics. The http tunnel and tcp tunnel parts of the command output below shows that WAN optimization has been processing HTTP and TCP packets.

```
get test wad 1
WAD manager process status: pid=113 n_workers=1 ndebug_workers=0
```

Enter the following command to display all test options:

```
get test wad

WAD process 82 test usage:
 1: display process status
 2: display total memory usage.
 99: restart all WAD processes
1000: List all WAD processes.
1001: display debug level name and values
1002: display status of WANOpt storages
1068: Enable debug for all WAD workers.
1069: Disable debug for all WAD workers.
2yxx: Set No. xx process of type y as diagnosis process.
 3: display all fix-sized advanced memory stats
 4: display all fix-sized advanced memory stats in details
500000..599999: cmem bucket stats (599999 for usage)
800..899: mem_diag commands (800 for help & usage)
800000..899999: mem_diag commands with 1 arg (800 for help & usage)
80000000..89999999: mem_diag commands with 2 args (800 for help & usage)
 60: show debug stats.
 61: discard all wad debug info that is currently pending
62xxx: set xxxM maximum output buffer size for WAD debug. 0, set back to default.
 68: Enable process debug
```

```

69: Disable process debug
98: gracefully stopping WAD process
9xx: Set xx workers(0: default based on user configuration.)

```

diagnose wad

Display diagnostic information about the WAN optimization daemon (wad).

```

diagnose wad console-log {disable | enable}
diagnose wad debug-url {disable | enable}
diagnose wad filter {clear | dport | dst | list | negate | protocol | sport | src | vd}
diagnose wad history {clear | list}
diagnose wad session {clear | list}
diagnose wad stats {cache | cifs | clear | crypto | ftp | http | list | mapi | mem |
    scan | scripts | summary | tcp | tunnel}
diagnose wad user {clear | list}
diagnose wad tunnel {clear | list}1
diagnose wad webcache {clear | list} {10min | hour | day | 30days}

```

Variable	Description
console-log	Enable or disable displaying WAN optimization log messages on the CLI console.
filter	Set a filter for listing WAN optimization daemon sessions or tunnels.
	clear reset or clear the current log filter settings.
	dport enter the destination port range to filter by.
	dst enter the destination address range to filter by.
history	list display the current log filter settings
	Display statistics for one or more WAN optimization protocols for a specified period of time (the last 10 minutes, hour, day or 30 days).
session	Display diagnostics for WAN optimization sessions or clear active sessions.
stats	Display statistics for various parts of WAN optimization such as cache statistics, CIFS statistics, MAPI statistics, HTTP statistics, tunnel statistics etc. You can also clear WAN optimization statistics and display a summary.
tunnel	Display diagnostic information for one or all active WAN optimization tunnels. Clear all active tunnels. Clear all active tunnels.
webcache	Display web cache activity for the specified time period.

Example diagnose wad tunnel list

Enter the following command to list all of the running WAN optimization tunnels and display information about each one. The command output shows 10 tunnels all created by peer-to-peer WAN optimization rules (auto-detect set to off).

```
diagnose wad tunnel list

Tunnel: id=100 type=manual
  vd=0 shared=no uses=0 state=3
  peer name=Web_servers id=100 ip=172.20.120.141
  SSL-secured-tunnel=no auth-grp=
  bytes_in=348 bytes_out=384

Tunnel: id=99 type=manual
  vd=0 shared=no uses=0 state=3
  peer name=Web_servers id=99 ip=172.20.120.141
  SSL-secured-tunnel=no auth-grp=
  bytes_in=348 bytes_out=384

Tunnel: id=98 type=manual
  vd=0 shared=no uses=0 state=3
  peer name=Web_servers id=98 ip=172.20.120.141
  SSL-secured-tunnel=no auth-grp=
  bytes_in=348 bytes_out=384

Tunnel: id=39 type=manual
  vd=0 shared=no uses=0 state=3
  peer name=Web_servers id=39 ip=172.20.120.141
  SSL-secured-tunnel=no auth-grp=
  bytes_in=1068 bytes_out=1104

Tunnel: id=7 type=manual
  vd=0 shared=no uses=0 state=3
  peer name=Web_servers id=7 ip=172.20.120.141
  SSL-secured-tunnel=no auth-grp=
  bytes_in=1228 bytes_out=1264

Tunnel: id=8 type=manual
  vd=0 shared=no uses=0 state=3
  peer name=Web_servers id=8 ip=172.20.120.141
  SSL-secured-tunnel=no auth-grp=
  bytes_in=1228 bytes_out=1264

Tunnel: id=5 type=manual
  vd=0 shared=no uses=0 state=3
  peer name=Web_servers id=5 ip=172.20.120.141
  SSL-secured-tunnel=no auth-grp=
  bytes_in=1228 bytes_out=1264

Tunnel: id=4 type=manual
  vd=0 shared=no uses=0 state=3
  peer name=Web_servers id=4 ip=172.20.120.141
  SSL-secured-tunnel=no auth-grp=
  bytes_in=1228 bytes_out=1264

Tunnel: id=1 type=manual
  vd=0 shared=no uses=0 state=3
  peer name=Web_servers id=1 ip=172.20.120.141
  SSL-secured-tunnel=no auth-grp=
  bytes_in=1228 bytes_out=1264

Tunnel: id=2 type=manual
```



```

vd=0 shared=no uses=0 state=3
peer name=Web_servers id=2 ip=172.20.120.141
SSL-secured-tunnel=no auth-grp=
bytes_in=1228 bytes_out=1264

Tunnels total=10 manual=10 auto=0

```

Example diagnose wad webcache list

This following command displays the web caching stats for the last 10 minutes of activity. The information displayed is divided into 20 slots and each slot contains stats for 30 seconds:

20 * 30 seconds = 600 seconds = 10 minutes

```

diagnose wad webcache list 10min
web cache history vd=0 period=last 10min

```

The first 20 slots are for HTTP requests in the last 10 minutes. Each slot of stats has four numbers, which is the total number of HTTP requests, the number of cacheable HTTP requests, the number of HTTP requests that are processed by the web cache (hits), and the number of HTTP requests that are processed without checking the web cache (bypass). There are many reasons that a HTTP request may bypass web cache.

total	cacheable	hits	bypass
-----	-----	-----	-----
36	10	3	1
128	92	1	10
168	97	2	3
79	56	0	3
106	64	5	3
180	118	6	11
88	53	7	3
80	43	4	4
107	44	9	2
84	12	0	2
228	139	52	10
32	2	0	5
191	88	13	7
135	25	40	3
48	10	0	8
193	13	7	7
67	31	1	2
109	35	24	6
117	36	10	5
22	0	0	4

The following slots are for video requests in the last 10 minutes. Each slot has two numbers for each 30 seconds: total number of video requests, and the number of video requests that are processing using cached data.

[illegible]

The following 20 slots are for traffic details in last 10 minutes. Each slot has four numbers for 30 seconds each.

--- LAN ---		--- WAN ---	
bytes_in	bytes_out	bytes_in	bytes_out
34360	150261	141086	32347
105408	861863	858501	100670
128359	1365919	1411849	127341
60103	602813	818075	59967
105867	1213192	1463736	97489
154961	1434784	1344911	158667
73967	370275	369847	70626
129327	602834	592399	123676
115719	663446	799445	111262
58151	724993	631721	59989
175681	2092925	1092556	166212
37805	33042	41528	37779
183686	1255118	1114646	172371
106125	904178	807152	81520
66147	473983	543507	66782
170451	1289530	1201639	165540
69196	544559	865370	68446
134142	579605	821430	132113
96895	668037	730633	89872
59576	248734	164002	59448

diagnose wad csvc

The `diagnose wad csvc` command refers to the `cache-service`. The next options to the command are listed in the table. Some will have there own sub options for refining the output or results.

Option	Description
<code>memory</code>	Cache service memory diagnostics

Option	Description
webcache	Webcache diagnostics
bytecache	Bytecache diagnostics
memcache	Memcache diagnostics
restart	Restart cache service

diagnose wad worker

The `diagnose wad worker` command has some settings that show useful WAD stats for one or all workers. The next options to the command are listed in the table. Some will have their own sub options for refining the output or results.

Option	Description
memory	WAD worker memory diagnostics.
tcp	TCP statistics.
ssl	SSL statistics.
tunnel	Tunnel statistics.
webcache	Webcache diagnostics.
bytecache	Bytecache diagnostics.
memcache	Memcache diagnostics.
restart	Restart workers.

diagnose wacs

Display diagnostic information for the web cache database daemon (wacs).

```
diagnose wacs clear
diagnose wacs reents
diagnose wacs restart
diagnose wacs stats
```

Variable	Description
<code>clear</code>	Remove all entries from the web cache database.
<code>reents</code>	Display recent web cache database activity.

Variable	Description
restart	Restart the web cache daemon and reset statistics.
stats	Display web cache statistics.

diagnose wadbd

Display diagnostic information for the WAN optimization database daemon (wadbd).

```
diagnose wadbd {check | clear | reents | restart | stats}
```

Variable	Description
check	Check WAN optimization database integrity.
clear	Remove all entries from the WAN optimization database.
reents	Display recent WAN optimization database activity.
restart	Restart the WAN optimization daemon and reset statistics.
stats	Display WAN optimization statistics.

diagnose debug application {wad | wccpd} [<debug_level>]

View or set the debug level for displaying WAN optimization and web cache-related daemon debug messages. Include a <debug_level> to change the debug level. Leave the <debug_level> out to display the current debug level. Default debug level is 0.

```
diagnose debug application wad [<debug_level>]
diagnose debug application wccpd [<debug_level>]
```

Variable	Description
wad	Set the debug level for the WAN optimization daemon.
wccpd	Set the debug level for the WCCP daemon.

diagnose test application wad 2200

The debug level 2200 switches the debug to explicit proxy mode. You have to enter this debug level first. After that you have to type the command again with a different debug level to check the different explicit proxy statistics. To list what each debug level shows, follow these steps in any FortiGate device:

1. Enable explicit proxy globally and in one interface, to start the wad process. If the wad process is *not* running, you *cannot* list the options.
2. Once the wad process starts, type:

```
diagnose test application wad 2200
diagnose test application wad //// Do not type any debug level value to list all the options.
```

This is the output you will get:

```
# diagnose test application wad 2200
Set diagnosis process: type=wanopt index=0 pid=114
# diagnose test application wad
WAD process 114 test usage:
1: display process status
2: display total memory usage
99: restart all WAD processes
1000: List all WAD processes
1001: display debug level name and values
1002: display status of WANOpt storages
1068: Enable debug for all WAD workers
1069: Disable debug for all WAD workers
2yxx: Set No. xx process of type y as diagnosis process
3: display all fix-sized advanced memory stats
4: display all fix-sized advanced memory stats in details
500000..599999: cmem bucket stats (599999 for usage)
800..899: mem_diag commands (800 for help & usage)
800000..899999: mem_diag commands with 1 arg (800 for help & usage)
80000000..89999999: mem_diag commands with 2 args (800 for help & usage)
60: show debug stats
61: discard all wad debug info that is currently pending
62xxx: set xxxM maximum output buffer size for WAD debug (0: set back to default)
68: Enable process debug
69: Disable process debug
98: gracefully stopping WAD process
20: display all listeners
21: display TCP port info
22: display SSL stats
23: flush SSL stats
24: display SSL mem stats
70: display av memory usage
71xxxx: set xxxMiB maximum AV memory (0: set back to default)
72: toggle av memory protection
73: toggle AV conserve mode (for debug purpose)
90: set to test disk failure
91: unset to test disk failure
92: trigger a disk failure event
100: display explicit proxy settings
101: display firewall policies
102: display security profile mapping for regular firewall policy
103: display Web proxy forwarding server and group
104: display DNS stats
105: display proxy redirection scan stats
106: list all used fqdns
107: list all firewall address
110: display current web proxy users
111: flush current web proxy users
112: display current web proxy user summary
113: display WAD fsso state
114: display HTTP digest stats
115: display URL patterns list of cache exemption or forward server
116: toggle dumping URL when daemon crashes
120: display Web Cache stats
```

```
121: flush Web Cache stats
122: flush idle Web cache objects
123: display web cache cache sessions
130: display ftpproxy stats
131: clear ftpproxy stats
132: list all current ftpproxy sessions
133: display all caughted webfilter profiles
200: display WANopt profiles
201: display all peers
202: display video cache rules (patterns)
203: display all ssl servers
210: toggle disk-based byte-cache
211: toggle memory-based byte-cache
212: toggle cifs read-ahead
221: display tunnel protocol stats
222: flush tunnel protocol stats
223: display http protocol stats
224: flush http protocol stats
225: display cifs protocol stats
226: flush cifs protocol stats
227: display ftp protocol stats
228: flush ftp protocol stats
229: display mapi protocol stats
230: flush mapi protocol stats
231: display tcp protocol stats
232: flush tcp protocol stats
233: display all protocols stats
234: flush all protocols stats
240: display WAD tunnel stats
241: display tunnel compressor state
242: flush tunnel compressor stats
243: display Byte Cache DB state
244: flush Byte Cache DB stats
245: display Web Cache DB state
246: flush Web Cache DB stats
247: display cache state
248: flush cache stats
249: display memory cache state
250: flush memory cache stats
261yxxx: set xxx concurrent Web Cache session for object storage y
262yxxx: set xxxK(32K, 64K,...) unconfirmed write/read size per Web Cache object for
        object storage y
263yxxxx: set xxxK maximum ouput buffer size for object storage y
264yxx: set lookup lowmark (only if more to define busy status) to be xx for object
        storage y
265yxxx: set xxxK maximum ouput buffer size for byte storage y
266yxxx: set number of buffered add requests to be xxx for byte storage y
267yxxxx: set number of buffered query requests to be xxxx for byte storage y
268yxxxxx: set number of concurrent query requests to be xxxxx for byte storage y
```

Chapter 7 - FortiView

- [Overview on page 924](#) outlines the role FortiView plays in FortiOS and its overall layout. This section also identifies which FortiGate platforms support the full FortiView features.
- [FortiView consoles on page 935](#) describes the various FortiView consoles available in FortiOS, including example scenarios, in most cases.
- [Reference on page 946](#) explains reference information for the various consoles in FortiView, and describes the assortment of filtering options, drilldown options, and columns available.
- [Troubleshooting FortiView on page 958](#) offers solutions to common technical issues experienced by FortiGate users regarding FortiView.

What's new in FortiOS 6.0

The following list contains new FortiView features added in FortiOS 6.0 and subsequent versions.

What's New in FortiOS 6.0

These features first appeared in FortiOS 6.0.

Increased visibility of interactivity in FortiView

FortiOS 6.0 has improved visual feedback on mouse-over of FortiView charts, to increase visibility of interactive elements. Cursors have also been updated to display possible interactions.

Automatic synchronization of log display location

In previous versions, log display location could differ between Log & Report and FortiView, which could result in empty log screens if the two were not synchronized. Now, both log viewers automatically pick the best available log device. A different log device can be manually selected.

As a result, the associated CLI command `log gui-display location` has been removed.

Improved log display consistency at high load

Previous versions could display inconsistent log data when using Drill Down charts and when navigating between different log tables (in both **Log & Report** and **FortiView**). The maximum number of records now varies based on length that logs are kept, relative to device model size. Record numbers are configurable in `config report setting`.

Log database queries used to collect **Top Sources** and **Top Destinations** data are significantly more efficient due to improved indexing speed.

Improved FortiView > Cloud Application

Previously, the **Cloud Application** view required that you enable deep inspection on a firewall policy to display any results. In 6.0, when applications are detected with the "Cloud" behavior category, it will be recorded and displayed in **Cloud Application** view, even when deep inspection is disabled.

New Destination > Owner views and FortiView features

A new Owner category has been added to **FortiView > Destination** visualizations. Used for tracking multiple devices across remote networked clients based on FortiClient data, the Owner view also appears in Topology pages and drilldown charts for Realtime, Historical Disk, and FortiAnalyzer logging data.

Retrieve FortiView data from FortiCloud

In 6.0, FortiView will retrieve data from FortiCloud for a wider range of pages, now supporting every page that is also supported by FortiAnalyzer.

Supported pages are: **Sources, Destinations, Applications, Cloud Applications, Web Sites, Threats, WiFi Clients, System Events, VPN, Endpoint Vulnerability, Policies, Interfaces, All Sessions, Physical Topology, Logical Topology, and FortiSandbox.**

Various GUI fixes

Various other small interface fixes were implemented in 6.0:

- Bubble chart color now matches threat level, if one is set.
- New icons added for various device types.
- Colors of icons, charts, and other visualizations have been updated for clarity and contrast.

Purpose

FortiView is a comprehensive monitoring system for your network that integrates real-time and historical data into a single view on your FortiGate. It can log and monitor threats to networks, filter data on multiple levels, keep track of administrative activity, and more.

FortiView allows you to use multiple filters within the consoles, enabling you to narrow your view to a specific time (up to 24 hours in the past), by user ID or local IP address, by application, and many more. For more on FortiView's filtering options, see [Filtering options on page 947](#)

FortiView can be used to investigate traffic activity, such as user uploads/downloads or videos watched on YouTube, on a network-wide, user group, and individual-user level, with information relayed in both text and visual format. FortiView makes it easy to get an actionable picture of your network's internet activity.

The degree to which information can be logged will depend on which FortiGate unit you have. For more information, see [Enabling FortiView on page 925](#).

Overview

This section provides an overview of FortiView, its interface, and options, including the following:

Enabling FortiView

By default, FortiView is enabled on FortiGate running FortiOS firmware version 5.2 and above. You will find the FortiView consoles in the main menu. However, certain options will not appear unless the FortiGate has **Disk Logging** enabled.

Only certain FortiGate models support Disk Logging. A complete list of FortiGate platforms that support Disk Logging is provided in the matrix below.

To enable Disk Logging

1. Go to **Log & Report > Log Settings** and select the checkbox next to **Disk**.
2. **Apply** the change.

To enable Disk Logging - CLI

```
config log disk setting
    set status enable
end
```

FortiView feature support - platform matrix

Note that the following table identifies three separate aspects of FortiView in FortiOS, which are explained in greater detail below:

- [Basic feature support](#)
- [Historical data](#)
- [Disk logging](#)

Platform	Basic Feature Support	Disk Logging	Historical Data *
FG/FWF-30D/E Series	✓		
FG/FWF-50E	✓		
FG/FWF-51E	✓	✓	1 hour
FG-52E	✓	✓	1 hour
FG/FWF-60D/E	✓		
FG-61E	✓	✓	1 hour
FG/FWF-70D Series	✓		
FG-80D	✓	✓	1 hour
FG-80E	✓		

Platform	Basic Feature Support	Disk Logging	Historical Data *
FG-81E	✓	✓	1 hour
FG/FWF-90D	✓	✓	1 hour
FG-90E	✓		
FG-91E	✓	✓	1 hour
FG/FWF-92D Series	✓	✓	1 hour
FG-100D	✓	✓	24 hours
FG-100E	✓		
FG-101E	✓	✓	24 hours
FG-200D	✓	✓	24 hours
FG-200E	✓		
FG-201E	✓	✓	24 hours
FG-300D	✓	✓	24 hours
FG-300E	✓		
FG-301E	✓	✓	24 hours
FG-400E	✓		
FG-500D	✓	✓	24 hours
FG-500E	✓		
FG-501E	✓	✓	24 hours
FG-600D	✓	✓	24 hours
FG-800D	✓	✓	24 hours
FG-900D	✓	✓	24 hours
FG-1000D	✓	✓	24 hours
FG-1200D	✓	✓	24 hours
FG-1500D	✓	✓	7 days

Platform	Basic Feature Support	Disk Logging	Historical Data *
FG-2000E	✓	✓	7 days
FG-2500E	✓	✓	7 days
FG-3000D	✓	✓	7 days
FG-3100D	✓	✓	7 days
FG-3200D	✓	✓	7 days
FG-3700D/DX	✓	✓	7 days
FG-3800D	✓	✓	7 days
FG-3810D	✓	✓	7 days
FG-3815D	✓	✓	7 days
FG-3960E	✓	✓	7 days
FG-3980E	✓	✓	7 days
FG-5001D	✓	✓	7 days

✓ = Default support.

* Refer to section on Historical Data below.

Basic feature support

FortiView's consoles give insight into your user's traffic, not merely showing which users are creating the most traffic, but what sort of traffic it is, when the traffic occurs, and what kind of threat the traffic may pose to the network.

FortiView basic feature support consists of the following consoles:

- [Sources](#)
- [Destinations](#)
- [Interfaces](#)
- [Policies](#)
- [All Sessions](#)
- [Applications](#)

The complete array of features in FortiView requires disk logging enabled (see below). It includes those consoles listed above as well as the following:

- [WiFi Clients](#)
- [Cloud Applications](#)
- [Web Sites](#)

- [Threats](#)
- [VPN](#)

Historical data

Not all consoles have the same available historical data options, depending on whether or not your traffic is locally stored.

Below is a table showing which features are available for units using local storage, including the historical data options.



Only FortiGate models 100D and above support the 24 hour historical data.

Features	With Local Storage				Without Local Storage			
	Now	5 min	1 hr	24 hr *	Now	5 min	1 hr	24 hr
Sources	✓	✓	✓	✓	✓			
Destinations	✓	✓	✓	✓	✓			
Interfaces	✓	✓	✓	✓				
Policies	✓	✓	✓	✓				
All Sessions	✓	✓	✓	✓	✓			
Applications	✓	✓	✓	✓	✓			
WiFi Clients		✓	✓	✓				
Cloud Applications	✓	✓	✓	✓	✓			
Web Sites	✓	✓	✓	✓				
Threats		✓	✓	✓				
Threat Map	✓				✓			
FortiSandbox		✓	✓	✓				
System Events		✓	✓	✓				
VPN		✓	✓	✓				

* Not available for desktop models with SSD.

7-day time display

As mentioned previously, certain models support 7-day time display. These models are listed below:

- FortiGate 1000D
- FortiGate 1500D
- FortiGate 3700DX
- FortiGate 3700D

The option for 7-day time display, however, can only be configured in the CLI using the following command:

```
config log setting
    set fortiview-weekly-data {enable|disable}
end
```

Disk logging

Only certain FortiGate models support Disk Logging (see above).

To enable Disk Logging, go to **Log & Report > Log Settings**, and select the checkbox next to **Disk** and apply the change. Some devices will require disk logging to be enabled in the CLI, using the following command:

```
config log disk setting
    set status enable
end
```

Configuration dependencies

Most FortiView consoles require the user to enable several features to produce data. The following table summarizes the dependencies:

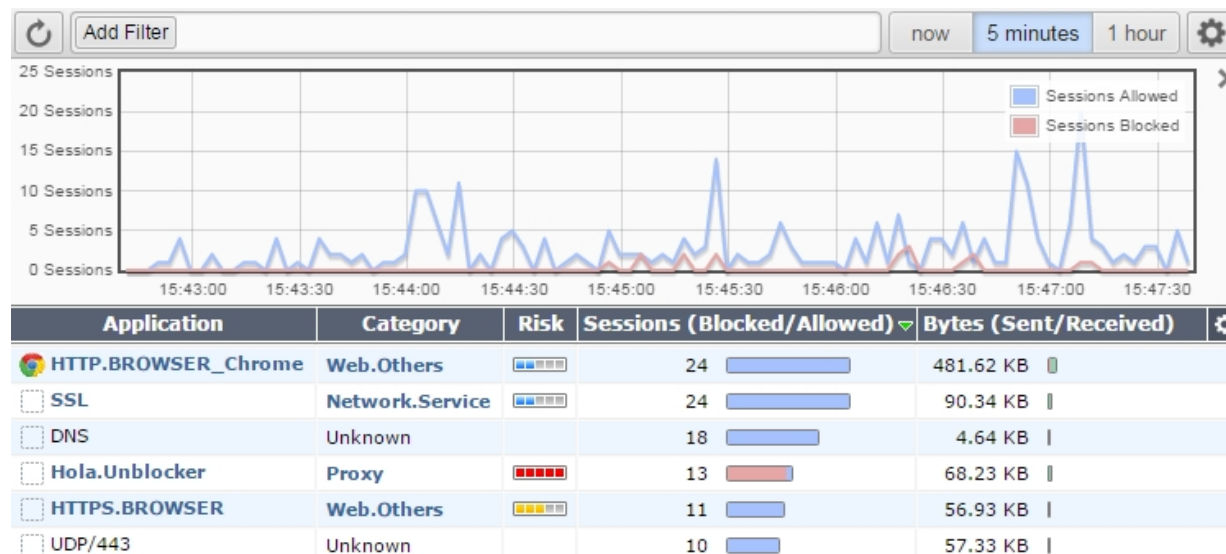
Feature	Dependencies (Realtime)	Dependencies (Historical)
Sources	None, always supported	Traffic logging enabled in policy
Destinations	None, always supported	Traffic logging enabled in policy
Interfaces	None, always supported	Disk logging enabled Traffic logging enabled in policy
Policies	None, always supported	Disk logging enabled Traffic logging enabled in policy
All Sessions	None, always supported	Traffic logging enabled in policy
Applications	None, always supported	Disk logging enabled Traffic logging enabled in policy Application control enabled in policy

Feature	Dependencies (Realtime)	Dependencies (Historical)
WiFi Clients	SSID must be in Tunnel mode	Disk logging enabled Traffic logging enabled in policy SSID must be in Tunnel mode
Cloud Applications	Not supported	Disk logging enabled Application control enabled in policy SSL "deep inspection" enabled in policy Deep application inspection enabled in application sensor Extended UTM log enabled in application sensor
Web Sites	Disk logging enabled Web Filter enabled in policy "web-url-log" option enabled in Web Filter profile	Disk logging enabled Web Filter enabled in policy "web-url-log" option enabled in Web Filter profile
Threats	Not supported	Disk logging enabled Traffic logging enabled in policy Threat weight detection enabled
Threat Map	None, always supported	Disk logging enabled Traffic logging enabled in policy Threat weight detection enabled
FortiSandbox	Not supported	Disk logging enabled Traffic logging enabled in policy
System Events	Not supported	Disk logging enabled
VPN	Not supported	Disk logging enabled Traffic logging enabled in policy

FortiView interface

FortiView lets you access information about the traffic activity on your FortiGate, visually and textually. FortiView is broken up into several consoles, each of which features a top menu bar and a graph window, as seen in the following image:

FortiView Application console sorted by Sessions (Blocked/Allowed)



The top menu bar features:

- a **Refresh** button, which updates the data displayed,
- a **Filter** button, for filtering the data by category,
- a **Settings** button (containing additional viewing settings and a link to the Threat Weight menu).
- a drop-down menu of different views:
 - **Time Display** (options: now, 5 minutes, 1 hour, or 24 hours),
 - **Table View**
 - **Timeline View**
 - **Bubble Chart**¹
 - **Country Map**

¹ For information on the Bubble Chart, refer to [Bubble chart visualization on page 932](#).



Certain views are only available in specific consoles.

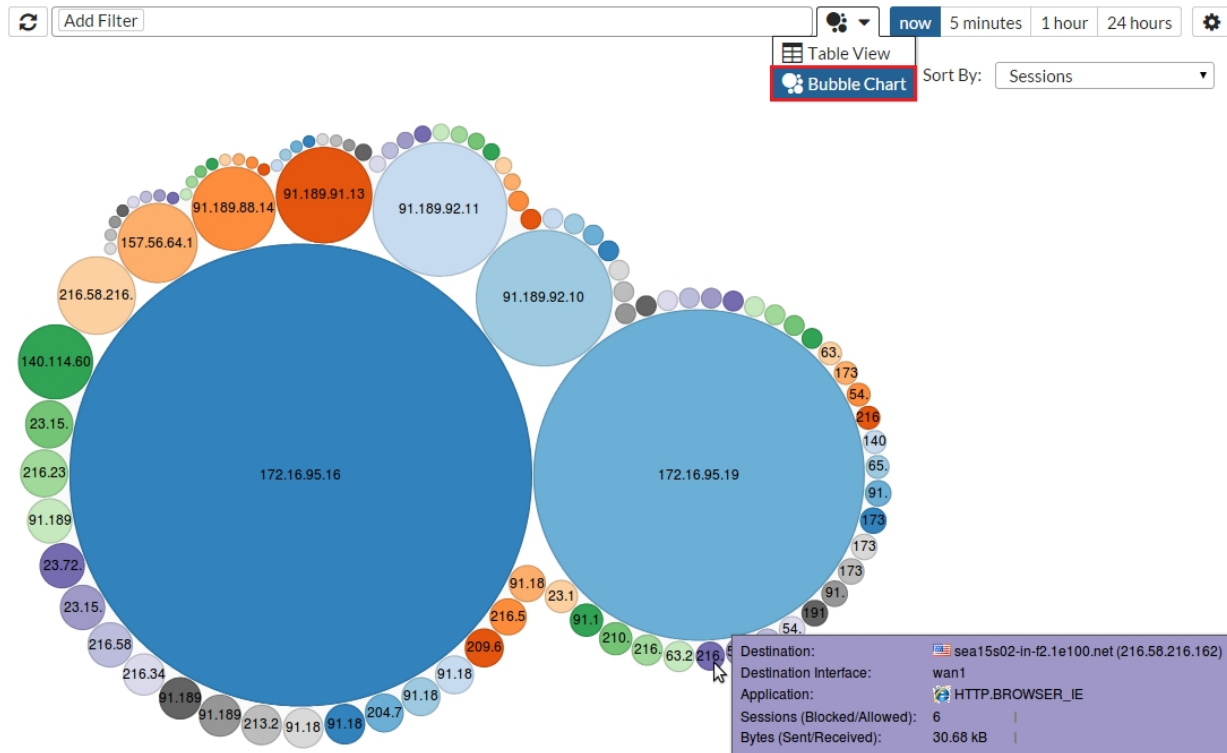
The FortiView graph

The graph window can be hidden using the **X** in the top right corner, and re-added by selecting **Show Graph**. To zoom in on a particular section of the graph, click and drag from one end of the desired section to the other. This will appear in the **Time Display** options as a **Custom** selection. The minimum selection size is 60 seconds.



Only FortiGate models 100D and above support the 24 hour historical data.

Bubble chart visualization



Notes about the Bubble chart:

- It is possible to sort on the Bubble chart using the **Sort By:** dropdown menu.
- The size of each bubble represents the related amount of data.
- Place your cursor over a bubble to display a tool-tip with detailed info on that item.
- You can click on a bubble to drilldown into greater (filtered) detail.

Links created between FortiView and View/Create Policy

The **Policy** column in FortiView consoles and the Log Viewer pages includes a link, which navigates to the IPv4 or IPv6 policy list and highlights the policy.

Right-clicking on a row in FortiView or the Log Viewer has menu items for **Block Source**, **Block Destination** and **Quarantine Source** where appropriate columns are available to determine these values. When multiple rows are selected, the user will be prompted to create a named **Address Group** to contain the new addresses.


When the user clicks **Block Source** or **Block Destination** they are taken to a policy creation page with enough information filled in to create a policy blocking the requested IP traffic.

The policy page will feature an informational message block at the top describing the actions that will be taken. Once the user submits the form, the requisite addresses, groups and policy will be created at once.

If the user clicks on **Quarantine User** then they will be prompted for a duration. They may also check a box for a **Permanent Ban**. The user can manage quarantined users under **Monitor > Quarantine Monitor**.

Visualization support for the Admin Logins page

A useful chart is generated for Admin login events under **FortiView > Admin Logins**. You can view the information in either **Table View** or **Timeline View** (shown below). In Timeline View, each line represents an administrator, with individual sessions indicated per administrator line. When you hover over a particular timeline, detailed information appears in a tooltip.



Realtime visualization

To enable realtime visualization:

1. Click on the **Settings** icon next to the upper right-hand corner and select **Auto update realtime visualizations**. An option is displayed to set the **Interval (seconds)**. The maximum value is 300.
2. Enter a desired **Interval** and click **Apply**.

Accelerated sessions

When viewing sessions in the [All Sessions](#) console, information pertaining to NP4/ NP6 acceleration is now reflected via an appropriate icon in the table. The tooltip for the icon includes the NP chip type and its total number of accelerated sessions.

Filtering on accelerated sessions

You can filter the console on 'FortiASIC' ('Accelerated' versus 'Not Accelerated') sessions.

WHOIS Lookup anchor for public IPv4 addresses

A Reverse IP lookup is possible using the WHOIS lookup icon available when you mouse over a public IP address in a FortiView log. If you left-click on the lookup icon, a new tab is opened in your browser for www.networksolutions.com, and a lookup is performed on the selected IP address (this option persists after drilling down one level in FortiView).

FortiView consoles

This section describes the following log filter consoles available in FortiView:

- [Sources on page 936](#) displays detailed information on the sources of traffic passing through the FortiGate, and the section covers how you can investigate an unusual spike in traffic to determine which user is responsible.
- [Destinations on page 936](#) displays detailed information on user destination-accessing through the use of drill down functionality.
- [Applications on page 937](#) displays Applications used on the network that have been recognized by Application Control, and this section shows how you can view what sort of applications individual employees are using.
- [Cloud Applications on page 937](#) displays Web/Cloud Applications used on the network, and this section shows how you can drill down to access detailed data on cloud application usage, e.g. YouTube.
- [Web Sites on page 938](#) displays websites visited as part of network traffic that have been recognized by Web Filtering, and this section shows how you can investigate instances of proxy avoidance, which is the act of circumventing blocks using proxies.
- [Threats on page 939](#) monitors threats to the network, both in terms of their Threat Score and Threat Level.
- [WiFi Clients on page 940](#) displays a list of all the devices connected to the WLAN.
- ["Traffic Shaping" on page 940](#) displays a list of existing Traffic Shapers, detailing their bandwidth use and which traffic is being shaped by each shaper.
- [System Events on page 940](#) displays security events detected by FortiOS, providing a name and description for the events, an assessment of the event's severity level, and the number of instances the events were detected.
- [VPN on page 941](#) displays how users can access information on any VPNs associated with their FortiGate.
- ["Endpoint Vulnerability" on page 942](#) displays a list of Vulnerability events detected by the FortiGate on networked devices, along with links to further vulnerability information and databases.
- [Threat Map on page 942](#) provides a geographical display of threats, in realtime, from international sources as they arrive at your FortiGate.
- [Policies on page 943](#) displays what policies are in affect on your network, what their source and destination interfaces are, how many sessions are in each policy, and what sort of traffic is occurring.
- [Interfaces on page 944](#) displays the number of interfaces connected to your network, how many sessions there are in each interface, and what sort of traffic is occurring.
- [FortiSandbox on page 944](#) displays FortiSandbox activity. FortiSandbox detects and analyzes advanced attacks designed to bypass traditional security defenses, and has a wide array of features that allow it to prevent future attacks from occurring again.
- [All Sessions on page 945](#) displays complete information on all FortiGate sessions, with the ability to filter sessions by port number and application type.

Sources

The **Sources** console provides information about the sources of traffic on your FortiGate unit.

This console can be filtered by Country, Destination Interface, Policy, Result, Source, and Source Interface. For more on filters, see [Filtering options](#).

Specific devices and time periods can be selected and drilled down for deep inspection.

Scenario: Investigating a spike in traffic

A system administrator notices a spike in traffic and wants to investigate it. From the **Sources** window, they can determine which user is responsible for the spike by following these steps:

1. Go to **FortiView > Sources**.
2. In the graph display, click and drag across the peak that represents the spike in traffic.
3. Sort the sources by bandwidth use by selecting the **Bytes (Sent/Received)** header.
4. Drill down into whichever source is associated with the highest amount of bandwidth use by double-clicking it. From this screen, you have an overview of that source's traffic activity.
5. Again, in either the **Applications** or **Destinations** view, select the **Bytes (Sent/Received)** header to sort by bandwidth use.
6. Double-click the top entry to drill down to the final inspection level, from which you can access further details on the application or destination, and/or apply a filter to prohibit or limit access.



Only FortiGate models 100D and above support the 24 hour historical data.

Destinations

The **Destinations** console provides information about the destination IP addresses of traffic on your FortiGate unit, as well as the application used. You can drill down the displayed information, and also select the device and time period, and apply search filters.

This console can be filtered by Country, Destination Interface, Destination IP, Policy, Result, and Source Interface. For more on filters, see [Filtering options](#).

Scenario: Monitoring destination data

The Destinations console can be used to access detailed information on user destination-accessing through the use of the console's drilldown functionality. In this scenario, the console is used to find out more about a particular user's Facebook usage patterns over a 24-hour period:

1. Go to **FortiView > Destinations**.
2. Select **1 hour** from the Time Display options at the top right corner of the console.
3. The easiest way to locate most destinations is to scan the Applications column for the name of the application. Once the session containing Facebook has been located, double-click it to access the Destination summary window.

4. Locate Facebook in the Applications column and double-click it to view the Facebook drilldown page. From here, detailed information regarding the user's Facebook session can be accessed.



Only FortiGate models 100D and above support the 24 hour historical data.

Applications

The **Applications** console provides information about the applications being used on your network.

This console can be filtered by Application, Country, Destination Interface, Policy, Result, and Source Interface. For more on filters, see [Filtering options](#).

Specific devices and time periods can be selected and drilled down for deep inspection.



In order for information to appear in the **Applications** console, Application Control must be enabled in a policy.

Scenario: Viewing application usage

A manager is interested in the office internet habits of their employees:

1. Go to **FortiView > Applications**, to view the list of applications accessed by the users on your network. Use the time-frame options to view what applications were used in those time periods (from now, 5 minutes, 1 hour, or 24 hours).
2. From **Sessions (Blocked/Allowed)** and **Bytes (Sent/Received)**, you can see how much traffic has been generated. Click these columns to show the traffic in descending order.
3. You notice that a social media application has created the most traffic of all the applications, and so it's at the top of the list. Drill down into the application by double-clicking or right-clicking and select **Drill Down to Details**.
4. You are directed to a summary page of the social media application. From here, you can see which specific user has made the most use of the application.



Only FortiGate models 100D and above support the 24 hour historical data.

Cloud Applications

The **Cloud Applications** console provides information about the cloud applications being used on your network. This includes information such as:

- The names of videos viewed on YouTube (visible by hovering the cursor over the session entry)
- Files uploaded and downloaded from cloud hosting services such as Dropbox
- Account names used for cloud services

Two different views are available for the Cloud Applications: **Applications** and **Users** (located in the top menu bar next to the time periods). **Applications** shows a list of the programs being used. **Users** shows information on the individual users of the cloud applications, including the username, if the FortiGate was able to view the login event.

This console can be filtered by Cloud Application and Result. For more on filters, see [Filtering options](#).



In order for information to appear in the **Cloud Applications** console, an application control profile (that has Deep Inspection of Cloud Applications turned on) must be enabled in a policy, and SSL Inspection must use `deep-inspection`.

Scenario: Viewing cloud application usage data

From the Cloud Applications console, users can drill down to access detailed data on cloud application usage data. In this scenario, the console is used to determine the network's most frequent user of YouTube over a 24-hour period, and find out more about their usage patterns.

1. Go to **FortiView > Cloud Applications**.
2. Select **Applications** view from the top menu bar if it is not already selected.
3. Select **24 Hours** from the Time Display options.
4. Find **YouTube** under the Application column and double-click it (or right-click and select **Drill down for details...**). This will open the YouTube stats window.
5. To determine the user who has accessed YouTube the most frequently, sort the column entries by **Sessions** by selecting the column header of the same name.
6. Double-click (or right-click and select **Drill down for details...**) the top-bandwidth YouTube user to view detailed stats, including the names of videos watched by the user and the date and time each video was accessed.



Only FortiGate models 100D and above support the 24 hour historical data.

Web Sites

The **Web Sites** console lists the top allowed and top blocked web sites. You can view information by domain or by FortiGuard categories by using the options in the top right corner. Each FortiGuard category can be selected in order to see a description of the category and several example sites, with content loaded from FortiGuard on demand.

This console can be filtered by Domain and Result. For more on filters, see [Filtering options](#).



In order for information to appear in the **Web Sites** console, web filtering must be enabled in a policy, with FortiGate Categories enabled.

Scenario: Investigating an instance of Proxy Avoidance

In this scenario, the Categories view will be used to investigate an instance of Proxy Avoidance, one of the Categories recognized by FortiOS. Proxy Avoidance denotes the use of a proxy site in order to access data that

might otherwise be blocked by the server.

1. Go to **FortiView > Web Sites** to open the Web Sites console.
2. Select **Categories** from the top bar menu to enter Categories view.
3. Scan the **Categories** column and locate the instance of Proxy Avoidance, then double-click it to enter its drilldown screen.



Only FortiGate models 100D and above support the 24 hour historical data.

Threats

The **Threats** console lists the top users involved in incidents, as well as information on the top threats to your network.

The following incidents are considered threats:

- Risk applications detected by application control
- Intrusion incidents detected by IPS
- Malicious web sites detected by web filtering
- Malware/botnets detected by antivirus

This console can be filtered by Country, Destination Interface, Policy, Result, Security Action, Source Interface, Threat, and Threat Type. For more on filters, see [Filtering options](#).



In order for information to appear in the **Threats** console, Threat Weight Tracking must be enabled.

Scenario: Monitoring Threats to the Network

Some users have high Threat Scores. The Threats console can be used to view all threats and discover why such high scores are being shown:

1. Go to **FortiView > Threats**. In the graph display, click and drag across the peak that represents the spike in threat score.
2. Sort the threats by score or level by selecting the **Threat Score (Blocked/Allowed)** or the **Threat Level** headers respectively.
3. You see that a specific threat's Threat Level is at Critical. Drill down into the threat by double-clicking or right-clicking and select **Drill down to details**.
4. From this summary page, you can view the source IPs and the number of sessions that came from this threat. Double-click on one of them.
5. The following page shows a variety of statistics, including **Reference**. The URL next to it will link you to a FortiGuard page where it will display the description, affected products, and recommended actions, if you are not familiar with the particular threat.



Only FortiGate models 100D and above support the 24 hour historical data.

WiFi Clients

The **WiFi Clients** console shows a list of all the devices connected to the WLAN. The type of device, source, number of sources blocked and allowed, and bytes sent and received are displayed. The source's Service Set Identifier (SSID) is also displayed in the **Source SSID** column. An SSID is a case sensitive, 32 character alphanumerical identifier that acts as a password when a mobile device tries to connect to the WLAN.

This console can be filtered by AP, Device Type, Result, Source Device, Source IP, Source SSID, and User. For more on filters, see [Filtering options](#).

Scenario: Determining the threat risk of an individual WiFi client

In this scenario, the administrator will use the WiFi Clients FortiView console to determine the risk levels associated with an individual WiFi client, and then drilldown into that client to determine where the risk originates and who might be the offending user/IP.

1. Go to **FortiView > WiFi Clients** and view the device list table.
2. Double-click on a device to filter on that source.
3. Under the **Risk** column, identify the items that present the greatest risk (using the **Applications**, **Destinations**, **Threats**, and/or **Sessions** tabs, for example).
4. Right-click these items for further action.

Traffic Shaping

The **Traffic Shaping** console provides information about FortiGate Traffic Shapers that are currently in effect. This console can be filtered by Traffic Shaper Name. For more on filters, see [Filtering options](#).

A number of columns available in FortiView are only available in Traffic Shaping. For example, the **Shaper** column displays the name of the Shaper, which can be used to monitor the traffic being shaped by Bytes Sent, Received, and Dropped, so that bandwidth patterns and Shaper effectiveness can be analyzed.



Only FortiGate models 100D and above support the 24 hour historical data.

System Events

The **System Events** console lists security events detected by FortiOS, providing a name and description for the events, an assessment of the event's severity level (**Alert**, **Critical**, **Emergency**, **Error**, or **Warning**), and the number of instances the events were detected.

Two other FortiView pages from 5.4 have been wrapped into the **System Events** page as of 5.6: **Admin Logins**, and **Failed Authentication**.

This console can be filtered by Event Name, Result, and Severity. For more on filters, see [Filtering options on page 947](#).

Scenario: Investigate network security events

System Events can be used in conjunction with All Sessions to see what network security events took place, and specifically see what action was taken upon their detection:

1. Go to **FortiView > System Events** to see what and how many network events have taken place, as well as how severe they are in terms of the threat they pose to the network.
2. You see that a particular event has warranted a severe rating, and has allowed traffic to bypass the firewall. Note when the event took place, and go to **FortiView > All Sessions**, to see more information pertaining to the security event.
3. From this console, you can determine the system event's source, how much traffic was sent and received, and the security action taken in response to this security event. These actions differ, depending upon the severity of the security event. See the entry for **Security Action** in [Columns displayed on page 950](#).



Only FortiGate models 100D and above support the 24 hour historical data.

VPN

From the **VPN** console, users can access information on any VPNs associated with their FortiGate. From the initial window, a list of all the associated VPNs is provided, along with general information, such as number of user connections and VPN type. By double-clicking on an individual VPN (or right-clicking and selecting **Drill down for details...**), users can access more specific data on that VPN.

Logs in the VPN console can be sorted by number of connections, last connection time, or data sent/received by selecting the column headers.

This console can be filtered by Result, User Name, and VPN Type. For more on filters, see [Filtering options on page 947](#).



Certain dashboard options will not appear unless your FortiGate has Disk Logging enabled.

Furthermore, only certain FortiGate models support Disk Logging — refer to the [FortiView feature support - platform matrix on page 925](#) for more information.

To enable Disk Logging, go to **Log & Report > Log Settings**, and select the checkbox next to **Disk** and apply the change.

Scenario: Investigating VPN user activity

The VPN console can be used to access detailed data on VPN-user activity via the use of the drill down windows. In this scenario, the administrator looks into the usage patterns of the IPsec user who has most frequently connected to the network.

1. Go to **FortiView > VPN** to view the VPN console.
2. Select the **Connections** column header to sort the entries by number of connections to the network.

3. Locate the top user whose VPN Type is **ipsec** and double-click the entry to enter that user's drill down screen.
4. To get the most representative data possible, sort the entries by bandwidth use by selecting the **Bytes (Sent/Received)** column header. Double-click the top entry to enter the drill down window for that connection instance.

From this screen, the administrator can find out more about the specific session, including the date/time of access, the XAuth (Extensible Authentication) User ID, the session's Tunnel ID, and more.



Only FortiGate models 100D and above support the 24 hour historical data.

Endpoint Vulnerability

The **Endpoint Vulnerability** console lists the top devices and vulnerabilities detected, organized either by frequency or risk level.

This console can be filtered by Vulnerability Name, Severity, Vulnerability Category, CVE-ID, or Host Count. For more on filters, see [Filtering options](#).

The Vulnerabilities detected by the FortiGate use definitions created by FortiGuard, and every vulnerability in FortiView contains a link to the respective FortiGuard Labs documentation page (under the '**Vulnerability ID**' column) and the Common Vulnerabilities and Exposures documentation page (under the '**CVE-ID**' column.)

Scenario: Monitoring Vulnerabilities on the Network

When a vulnerability appears in log data, you can use the FortiView page to see more information about it. The Endpoint Vulnerability console can be used to view and track all historical vulnerabilities:

1. Go to **FortiView > Endpoint Vulnerability**. In the upper right, select **Vulnerability**.
2. Sort the threats by frequency by selecting the **Host Count** header.
3. You see that a frequent vulnerability's Severity is at Critical. Drill down into the threat by double-clicking or right-clicking and select **Drill down to details**.
4. From this summary page, you can view the source IPs and devices on which this vulnerability was detected, and also the **Scan Time**. Double-click on one of them.
5. The chart will be filtered to display the specific Endpoint and Vulnerability, offering more granular data about the vulnerability, including its Category and the FortiClient ID of the device. You can access the CVE and FortiGuard links from this page to learn more.



Only FortiGate models 100D and above support the 24 hour historical data.

Threat Map

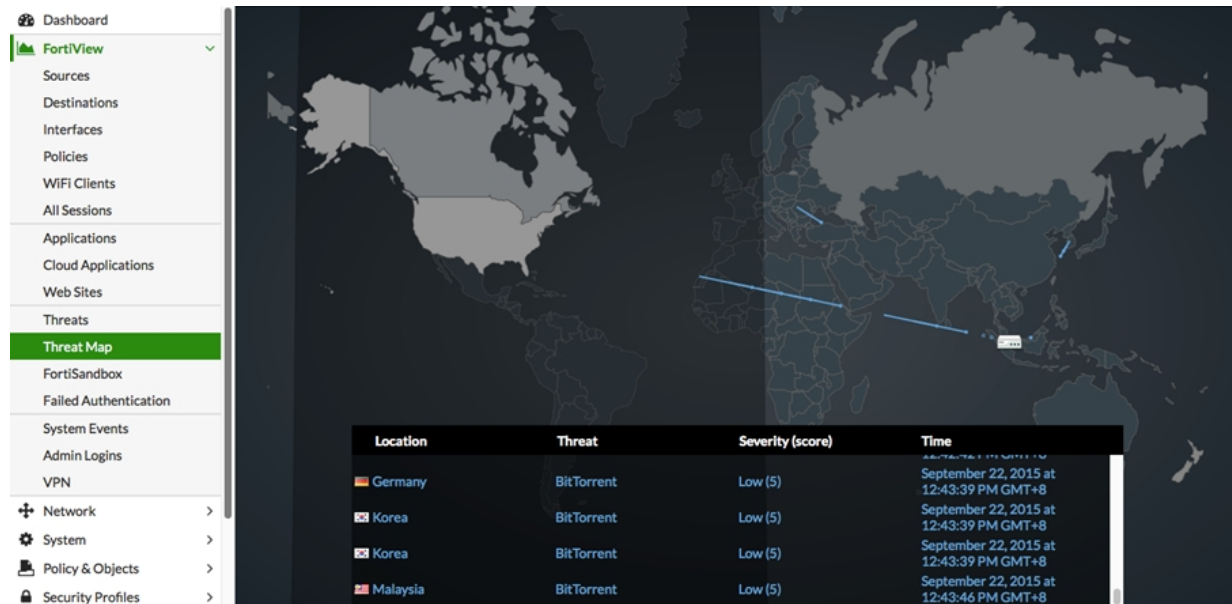
The **Threat Map** console displays network activity by geographic region. Threats from various international destinations will be shown, but only those arriving at your destination, as depicted by the FortiGate. You can place your cursor over the FortiGate's location to display the device name, IP address, and the city name/location.

A visual lists of threats is shown at the bottom, displaying the location, severity, and nature of the attacks. The color gradient of the darts on the map indicate the traffic risk, where red indicates the more critical risk.

Unlike other FortiView consoles, this console has no filtering options, however you can click on any country to drill down into greater (filtered) detail.



Only FortiGate models 100D and above support the 24 hour historical data.



Scenario: Investigate various international threats

The Threat Map console can be used to regionalize areas that you are more interested in, and disregard regions that you are not interested in:

1. Go to **FortiView > Threat Map** to see a real-time map of the globe. This will show various incoming threats from multiple destinations around the world, depending upon where the FortiGate is placed on the map.
2. You are not interested with threats that are being sent to Eastern Europe, however you are concerned with threats that may be sent to a city in North America. Click and drag the FortiGate to the approximate location where you would like to monitor the incoming threats.
3. To see which countries are sending the more severe threats to your region/location, either see where the red darts are coming from, or check the visual lists of threats at the bottom.

Policies

The **Policies** console shows what policies are in affect on your network, what their source and destination interfaces are, how many sessions are in each policy, and what sort of traffic is occurring, represented in bytes sent and received.

This console can be filtered by Country, Destination Interface, Destination IP, Policy, Source, Source Device, and Source Interface. For more on filters, see [Filtering options](#).



Only FortiGate models 100D and above support the 24 hour historical data.

Scenario: Investigate which policies are in effect

You can click on policy IDs to drill down to the policy list and see what policy's are in effect for specific interfaces, how many sessions have occurred, how many of those with the policy have been blocked, and more:

1. Go to **FortiView > Policies**, and double-click on a policy ID to drill down.
2. You will be redirected to a summary screen of the policy ID. From here you can view the source IP of where the policy has been used, what source interface has been using the particular policy, and to verify what sort of threat scores have been measured, both blocked and allowed.

Interfaces

The **Interfaces** console lists the total number of interfaces connected to your network, how many sessions there are in each interface, and what sort of traffic is occurring, represented in both bytes sent and received, and the total bandwidth used.

This console can be filtered by Country, Destination Interface, Destination IP, Policy, Result, Source, and Source Interface. For more on filters, see [Filtering options](#).



Only FortiGate models 100D and above support the 24 hour historical data.

Scenario: Investigate traffic spikes per user

The wan1 interface is showing a higher amount of traffic than usual. A system administrator uses the console to inspect which user (as represented by an IP address) is creating the spike in traffic:

1. Go to **FortiView > Interfaces** and double-click on wan1, or right click and select **Drill Down to Details...**
2. The console will drill down to a summary page of wan1, showing how many bytes are being sent and received, how much bandwidth is being used, and how many sessions are currently using this interface. You see the IP address of the user that is showing the most amount of traffic under **Source**.
3. You can further drill down to see the IP destination, the device, and the applications being used, and other options.

FortiSandbox

The **FortiSandbox** console detects and analyzes advanced attacks designed to bypass traditional security defenses, and has a wide array of features that allow it to prevent future attacks from occurring again.

This console can be filtered by Checksum, File Name, Source, Status, and User Name. For more on filters, see [Filtering options](#).



Only FortiGate models 100D and above support the 24 hour historical data.

All Sessions

The **All Sessions** console provides information about all FortiGate traffic. This console can be filtered by Application, Country, Destination Interface, Destination IP, Destination Port, NAT Source IP, NAT Source Port, Policy, Protocol, Source, Source Interface, Source IP, and Source Port. For more on filters, see [Filtering options](#).

This console has the greatest number of column options to choose from. To choose which columns you wish to view, select the column settings cog at the far right of the columns and select your desired columns. They can then be clicked and dragged in the order that you wish them to appear.

A number of columns available in FortiView are only available in All Sessions. For example, the **Action** column displays the type of response taken to a security event. This function can be used to review what sort of threats were detected, whether the connection was reset due to the detection of a possible threat, and so on. This would be useful to display alongside other columns such as the **Source**, **Destination**, and **Bytes (Sent/Received)** columns, as patterns or inconsistencies can be analyzed.

Similarly, there are a number of filters that are only available in All Sessions, one of which is **Protocol**. This allows you to display the protocol type associated with the selected session, e.g. TCP, FTP, HTTP, HTTPS, and so on.

Scenario: Filtering sessions by port number and application type

From the **All Sessions** console, a wide variety of filters can be applied to sort the session data. In this example, the All Sessions filters will be used to locate a specific user's recent Skype activity.

1. Go to **FortiView > All Sessions**.
2. Select **now** from the **Time Display** options if it is not already selected.
3. Select the **Filter** button, then select **Applications**. This will open a drop-down menu listing the applications that appear in the master session list. From this list, locate and select **Skype**, or type "Skype" into the Search Bar and hit **Enter**. This will filter the session list to only feature Skype usage.
4. Select the **Filter** button again, then select **Destination Port** from the drop-down menu, then locate and select the desired port number. This will add a second filter which will restrict the results to presenting only the Skype data associated with that port number.



Only FortiGate models 100D and above support the 24 hour historical data.

Reference

This section consists of reference information for the various consoles in FortiView. Each console has an assortment of filtering options, drilldown options, and columns that can be displayed. Since many of these options and columns persist through each console, the entire list of options and their descriptions is included below. Attempts have been made to identify the instances where an option or column is only available to a particular console.

This section includes:

Filtering options

When you select the **Add Filter** button, a drop-down list appears with a list of available filtering options. Available options differ based on which console is currently being viewed. The following table explains all of the available filtering options:

Filter option	Description
Accelerated Sessions	You can filter the console on 'FortiASIC' ('Accelerated' versus 'Not Accelerated') sessions.
AP	Filter by Access Point (AP) identification number.
Application	Filter by application name.
Checksum	Filter by checksum value. Checksums are reference digits used to represent the correct datasum of a packet in order to detect errors.
Cloud Application	Filter by cloud application name. Note: This filter is only available in the Cloud Applications console.
Country	Filter by the country from which the source accessed the server.
Destination Interface	Filter by the interface type used by the destination user, e.g. wan1.
Destination IP	Filter by the IP address used by the destination.
Destination Port	Filter by the port used by the destination. Note: This filter is only available in the All Sessions console, (viewing the now time display).
Domain	Filter by domain name. Note: This filter is only available in the Web Sites console.
Event Name	Filter by security event name. Note: This filter is only available in the System Events console.
File Name	Filter by file name. Note: This filter is only available in the FortiSandbox console.
Login Type	Filter by type of login (eg. WEP) associated with the displayed authentication attempt. Note: This filter is only available in the Failed Authentications console.

Filter option	Description
NAT Source IP	Filter by the NAT-translated source IP address. Note: This filter is only available in the All Sessions console, (viewing the now time display).
NAT Source Port	Filter by the NAT-translated source interface. Note: This filter is only available in the All Sessions console, (viewing the now time display).
Policy	Filter by the policy identification number.
Protocol	Filter by the protocol used by the source, e.g. tcp or udp. Note: This filter is only available in the All Sessions console, (viewing the now time display).
Result	Filter by the result of whatever security action was taken by FortiOs in the selected session, eg. Accept (all).
Security Action	Filter by the type of response taken to the security event. The types of possible actions are as follows: Allowed: No threat was detected and the connection was let through. Blocked: A threat was detected and the connection was not let through. Reset: A possible issue was detected and the connection was reset. Traffic Shape: Some data packets may have been delayed to improve system-wide performance.
Severity	Filter by the severity level (Critical , High , Medium or Low) associated with a security event.
Source	Filter by the source IP address.
Source IP	
Source Device	Filter by source device type, e.g. mobile.
Source Interface	Filter by the interface type used by the source user, e.g. wan1.
Source Port	Filter by the source interface. Note: This filter is only available in the All Sessions console, (viewing the now time display).

Filter option	Description
Source SSID	Filter by the Service Set Identifier (SSID) associated with the selected user. An SSID is a case sensitive, 32 character alphanumeric identifier that acts as a password attributed to a mobile device.
Status	Filter by the maliciousness of a file. The types of possible status' are Malicious, High, Medium, Low, Clean, Unknown, and Pending . Note: This filter is only available in the FortiSandbox console.
Threat	Filter by threat name and/or URL
Threat Type	Filter by threat category, e.g. <i>Illegal/Unethical</i> or <i>P2P</i> .
Type	Note: This filter is only available in the Failed Authentications console.
User Name	Filter by user name.
VPN Type	Filter by Virtual Private Network (VPN) protocol type, eg. <i>PPTP</i> . Note: This filter is only available in the VPN console.

Drill-Down Options

Double-click, or right-click, on any entry in a FortiView console and select **Drill Down to Details**, to view the following columns (options vary depending on the console selected):



Drill down options are available for all FortiView consoles except **All Sessions**, **Logical Topology**, and **Physical Topology**.

Option	Description
Applications	Select to drill down by application to view application-related information, including the application name, sessions blocked and allowed, bytes sent and received, and the risk level. You can sort entries by selecting the column header.
Sources	Select to drill down by rows to view source-related information, including IP address, device type, interface type, threat score, number of sessions blocked/allowed, and bytes sent/received. You can sort entries by selecting the column header.
Destinations	Select to drill down by destination to view destination-related information, including the IP address and geographic region, interface, threat score, number of sessions blocked and allowed, and bytes sent and received. You can sort entries by selecting the column header.

Option	Description
Countries	Select to drill down by country, including the number of sessions, bytes sent and received, and the bandwidth used. You can sort entries by selecting the column header.
Policies	Select to drill down by the policies in use, including source interface, destination interface, bytes sent and received, and bandwidth used. You can sort entries by selecting the column header.
Source Interfaces	Select to drill down by source interface, including bytes sent and received, and bandwidth used. You can sort entries by selecting the column header.
Destination Interfaces	Select to drill down by destination interface, including bytes sent and received, and bandwidth used. You can sort entries by selecting the column header.
Threats	Select to drill down by threat to view threat-related information, including the threat name, category, threat level, threat score, and number of sessions blocked and allowed. You can sort entries by selecting the column header.
Domains	Select to drill down by domain to view domain-related information, including domain name, category, browsing time, threat weight, number of sessions blocked/allowed, and bytes sent/received. You can sort entries by selecting the column header.
Categories	Select to drill down by category to view category-related information, including category name, browsing time, threat score, number of sessions blocked/allowed, and bytes sent/received. You can sort entries by selecting the column header.
Sessions	Select to drill down by sessions to view session-related information, including date/time, source, destination IP address and geographic region, application name, security action, security event, and bytes sent/received. You can sort entries by selecting the column header.

Columns displayed

The following columns appear in the initial window of the dashboards. Some columns may only be visible by selecting them from the column drop-down menu. Options vary depending on the dashboard selected.

Column name	Description
Action	<p>Displays the type of response taken to a security event. The types of possible actions are as follows:</p> <ul style="list-style-type: none"> • Allowed: No threat was detected and the connection was let through. • Blocked: A threat was detected and the connection was not let through. • Reset: A possible issue was detected and the connection was reset. • Traffic Shape: Some data packets may have been delayed to improve system-wide performance. <p>Note: This column is only available in the All Sessions console.</p>
Application	<p>Displays the application name and service. When Time Display is set to now, you can access further information about an application by selecting the column entry.</p>
Application Category	<p>Displays the type of application used in the selected session, e.g. video player, social media.</p> <p>Note: This column is only available in the All Sessions console.</p>
Application ID	<p>Displays the identification number associated with the application used in the selected session.</p> <p>Note: This column is only available in the All Sessions console.</p>
Application Risk Risk	<p>Displays the application risk level. You can hover the mouse cursor over the entry in the column for additional information, and select the column header to sort entries by level of risk.</p> <p>Risk uses a 5-point risk rating. The rating system is as follows:</p> <ul style="list-style-type: none"> • Critical: Applications that are used to conceal activity to evade detection. • High: Applications that can cause data leakage, are prone to vulnerabilities, or may download malware. • Medium: Applications that can be misused. • Elevated: Applications that are used for personal communications or can lower productivity. • Low: Business-related applications or other harmless applications.
Bandwidth	<p>Displays information for bandwidth calculated on a per-session level, providing administrators the ability to sort realtime bandwidth usage in descending order.</p>

Column name	Description
Browsing Time	<p>Displays the amount of time a user has spent browsing a web site (in seconds).</p> <p>Note: This column is only available in the Web Sites console, in Categories view..</p>
Bytes (Sent/Received)	<p>Displays the size of sent and received data packets, as measured in bytes. Select the column header to sort the entries by size.</p> <p>Note: This information is available on some consoles as two separate columns: Sent and Received.</p>
Category	Displays the category descriptor appropriate to whatever console is being displayed. For example, threat categories are displayed in the Threats console.
Cloud User	<p>Displays the users accessing cloud applications by IP address.</p> <p>Note: This column is only available in the Cloud Applications console, in Users view.</p>
Configuration Changes	<p>Displays the number of configuration changes made by the user. You can hover the mouse cursor over an entry for additional information.</p> <p>Note: This column is only available in the Admin Logins console.</p>
Connections	<p>Displays the number of VPN connections made by the selected user..</p> <p>Note: This column is only available in the VPN console.</p>
Country	<p>Displays the country from which the selected traffic is originating.</p> <p>Note: This column is only available in the Countries console.</p>
Destination	Displays the destination name, IP address and geographic region.
Destination Country	<p>Displays the country session data is being sent to.</p> <p>Note: This column is only available in the All Sessions console.</p>
Destination Interface	Displays which interface session data is being sent through, e.g. wan1.
Destination Port	<p>Displays the port number of the destination server being used to accept data.</p> <p>Note: This column is only available in the All Sessions console.</p>
Device	Displays the device IP address or Fully Qualified Domain Name (FQDN).

Column name	Description
Domain	Displays the domain associated with the selected web site, e.g. google.com. Note: This column is only available in the Web Sites console.
DST Nat IP NAT Destination	Displays the Network Address Translation (NAT) IP address associated with the destination server. Note: This column is only available in the All Sessions console.
DST Nat Port NAT Destination Port	Displays the Network Address Translation (NAT) port number associated with the destination server. Note: This column is only available in the All Sessions console.
Duration	Displays the amount of time (in seconds) a user has been logged in. Note: This column is only available in the Admin Logins console.
Event Name (Description)	Displays the name and description of the selected security event. Note: This column is only available in the System Events console.
Events	Displays the number of security events that occurred within a selected session. Note: This column is only available in the System Events console.
Expires	Displays the amount of time a session has (in seconds) before it is set to expire. Note: This column is only available in the All Sessions console, in now Time Display view.
Failed Logins	Displays the number of failed login attempts made by an administrator over the specified time period. Note: This column is only available in the Admin Logins console.
Files (Up/Down)	Displays the number of files uploaded and downloaded. Hover the mouse cursor over the entry in the column for additional information. Note: This column is only available in the Cloud Applications console.

Column name	Description
FortiASIC	<p>Displays the type of FortiASIC hardware acceleration used in the specified session, if present.</p> <p>Note: This column is only available in the All Sessions console, in the now Time Display view.</p>
Group	<p>Displays the group ID associated with the selected session.</p> <p>Note: This column is only available in the All Sessions console.</p>
Last Connection Time	<p>Displays the most recent instance of connection to the selected Virtual Private Network (VPN).</p> <p>Note: This column is only available in the VPN console.</p>
Level Threat Level	<p>Displays the threat level. Select the column header to sort entries by threat level.</p>
Log ID	<p>Displays the identification number for the data log associated with this entry.</p> <p>Note: This column is only available in the All Sessions console.</p>
Login IDs	<p>Displays the number of login IDs associated with the selected cloud application.</p> <p>Note: This column is only available in the Cloud Applications console, in Applications view.</p>
Login Type	<p>Displays the type of login (eg. WEP) associated with the displayed authentication attempt.</p> <p>Note: This column is only available in the Failed Authentications console.</p>
Logins	<p>Displays the number of successful logins made by an administrator over the specified time period.</p> <p>Note: This column is only available in the Admin Logins console.</p>
Pending	<p>Note: This column is only available in the FortiSandbox column, in Source view.</p>
Policy ID	<p>Displays the identification number of the policy under which the selected connection was allowed.</p>

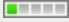




Column name	Description
Security Action	<p>Displays the action taken in response to the selected security event. The types of possible actions are as follows:</p> <ul style="list-style-type: none"> • Allowed: No threat was detected and the connection was let through. • Blocked: A threat was detected and the connection was not let through. • Reset: A possible issue was detected and the connection was reset. • Traffic Shape: Some data packets may have been delayed to improve system-wide performance.
Sessions	<p>Displays the number of sessions associated with the selected destination.</p> <p>Note: This column only appears in the Destinations console, in the now Time Display view.</p>
Sessions (Blocked/Allowed)	<p>Displays the number of sessions blocked and allowed by FortiOs.</p> <p>In some consoles, entries can be sorted by number of sessions by selecting the column header..</p>
Severity	Displays the severity level (Critical , High , Medium or Low) associated with the selected security event.
Source	Displays the source IP address and/or user ID, if applicable.
Source Interface	Displays which interface is being used by the destination server (eg. wan1).
Source Port	Displays the port number being used by the source server to send data.
Source SSID	<p>Displays the Service Set Identifier (SSID) associated with the selected user.</p> <p>Note: This column is only available in the Wifi Clients console.</p>
Src NAT IP NAT Source	Displays the Network Address Translation (NAT) IP address associated with the source server.
Src NAT Port NAT Source Port	Displays the Network Address Translation (NAT) port number associated with the source server.

Column name	Description
Status	<p>The types of possible status' are Malicious, High, Medium, Low, Clean, Unknown, and Pending.</p> <p>Note: This console is only available in the FortiSandbox console, in Files view.</p>
Submitted	<p>Displays the number of files submitted to the FortiSandbox for assessment in the selected session.</p> <p>Note: This column is only available in the FortiSandbox console, in Files view.</p>
Threat	Displays the threat type detected in the selected session.
Threat Score (Blocked/Allowed)	Displays the threat score value, a measurement of the total number of threats detected over the course of the session. You can select the column header to sort entries by threat score.
Threat Weight	Displays the threat weight profile associated with the selected session.
Timestamp	Displays the selected session's PHP timestamp.
User	
User Name	Displays the user name associated with the selected administrator.
Videos Played	<p>Displays the number of videos played via cloud applications.</p> <p>Note: This column is only available in the Cloud Applications console.</p>

Risk level indicators

There are currently two consoles within FortiView that display the Risk associated with the console: **Applications** and **Cloud Applications**. Each application pose different levels of risk to the network, represented by a colour code.

The following table identifies each risk level, from least to most severe:

Indicator	Risk	Description
	Green: <i>Risk Level 1</i>	<p>These applications have little to no risk level, with no assigned risk definition. Application file-sharing may result in data leakage, which would be a typical example of a low level risk.</p> <p>An example application would be the Google toolbar, or Dropbox.</p>
	Blue: <i>Risk Level 2</i>	<p>These applications have an elevated risk level and typically use excessive bandwidth. High bandwidth consumption can lead to increased operational costs.</p> <p>An example application would be Bittorrent.</p>
	Yellow: <i>Risk Level 3</i>	<p>These applications have a low risk level and are typically evasive.</p> <p>Evasive applications can lead to compliance risks, and could include applications such as JustinTV and GlypeProxy.</p>
	Orange: <i>Risk Level 4</i>	<p>These applications have a high risk level, and are defined as using both excessive and evasive bandwidth.</p> <p>Example applications would be AutoHideIP and PandoraTV.</p>
	Red: <i>Risk Level 5</i>	<p>Applications that have a high risk level are prone to malware or vulnerabilities that can introduce business continuity risks.</p>

Troubleshooting FortiView

No logging data is displayed

In order for information to appear in the FortiView consoles, disk logging must be selected for the FortiGate unit. To select disk logging, go to **Log & Report > Log Settings**.

Disk logging is disabled by default for some FortiGate units. To enable disk logging, enter the following command in the CLI:

```
config log disk setting
    set status enable
end
```

Only certain FortiGate models support Disk Logging — refer to the [FortiView feature support - platform matrix on page 925](#) for more information.

Logging is enabled, but data is not appearing

Some FortiView consoles require certain features to be enabled and working before they will display any data. For example, the Web Filtering FortiView page requires that a Web Filtering profile be configured in **Security Profiles > Web Filter** and then applied to a policy in **Policy & Objects > IPv4 Policy**.

First, ensure the feature is enabled in **System > Feature Visibility**, and then go to the appropriate page to make sure that the feature is being implemented. If it is working but is producing no data, FortiView will have nothing to display.

Chapter 8 - Fortinet Communication Ports and Protocols

What's new in FortiOS 6.0

The following list contains new communication ports and protocols features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- ["FortiOS DHCP options and auto DNS hostname for FortiManager details" on page 1011](#)

Introduction

This document contains a series of diagrams and tables showing the communication ports and protocols used between various Fortinet products:

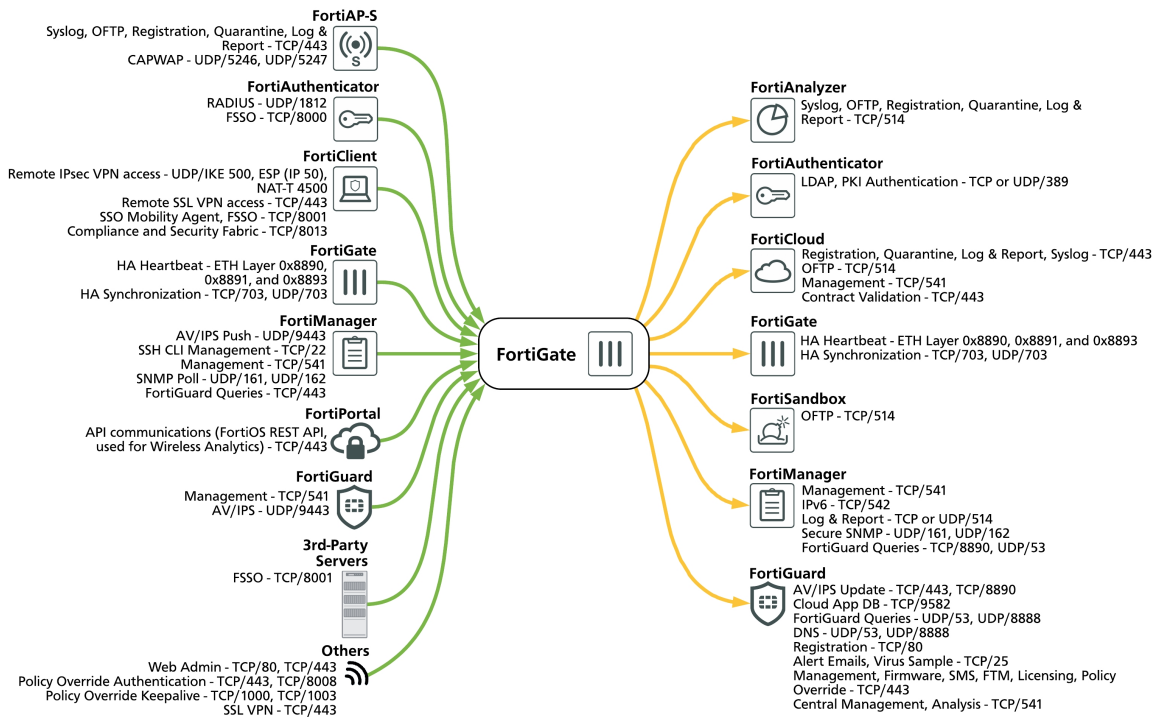
- [FortiGate](#)
- [FortiAnalyzer](#)
- [FortiAP-S](#)
- [FortiAuthenticator](#)
- [FortiClient](#)
- [FortiCloud](#)
- [FortiDB](#)
- [FortiGuard](#)
- [FortiMail](#)
- [FortiManager](#)
- [FortiPortal](#)
- [FortiSandbox](#)
- and [3rd-party servers](#) using FSSO.

Additionally, Fortinet's proprietary protocols are documented, showing what Fortinet products they operate with, how they behave, and how they carry out their roles:

- [FGCP - FortiGate Clustering Protocol](#)
- [FGSP - FortiGate Session Life Support Protocol](#)
- [FGFM - FortiGate to FortiManager Protocol](#)
- [SLBC - Session-aware Load Balancing Cluster](#)
- [Fortinet Security Fabric](#)
- [FortiTelemetry/On-Net/FortiClient Endpoint Compliance](#)
- [FortiGuard](#)
- [FortiLink](#)
- [FortiOS WAN optimization](#)
- [FSSO - Fortinet Single Sign-On](#)
- [OFTP - Optimized Fabric Transfer Protocol](#)
- [FortiClient EMS - Enterprise Management Server](#)

Some protocols contain CLI syntax that control their ports and functionality.

FortiGate open ports



Incoming Ports

Purpose	Protocol/Port
---------	---------------

FortiAP-S	Syslog, OFTP, Registration, Quarantine, Log & Report	TCP/443
	CAPWAP	UDP/5246, UDP/5247
FortiAuthenticator	RADIUS	UDP/1812
	FSSO	TCP/8000

Incoming Ports		
Purpose		Protocol/Port
FortiClient	Remote IPsec VPN access	UDP/IKE 500, ESP (IP 50), NAT-T 4500
	Remote SSL VPN access	TCP/443
	SSO Mobility Agent, FSSO	TCP/8001
	Compliance and Security Fabric	TCP/8013 (by default; this port can be customized)
FortiGate	HA Heartbeat	ETH Layer 0x8890, 0x8891, and 0x8893
	HA Synchronization	TCP/703, UDP/703
FortiGuard	Management	TCP/541
	AV/IPS	UDP/9443
FortiManager	AV/IPS Push	UDP/9443
	SSH CLI Management	TCP/22
	Management	TCP/541
	SNMP Poll	UDP/161, UDP/162
	FortiGuard Queries	TCP/443
	API communications (FortiOS REST API, used for Wireless Analytics)	TCP/443
Others	Web Admin	TCP/80, TCP/443
	Policy Override Authentication	TCP/443, TCP/8008
	Policy Override Keepalive	TCP/1000, TCP/1003
	SSL VPN	TCP/443
3rd-Party Servers	FSSO	TCP/8001 (by default; this port can be customized)

Outgoing Ports		
Purpose		Protocol/Port
FortiAnalyzer	Syslog, OFTP, Registration, Quarantine, Log & Report	TCP/514
FortiAuthenticator	LDAP, PKI Authentication	TCP or UDP/389
FortiCloud	Registration, Quarantine, Log & Report, Syslog	TCP/443
	OFTP	TCP/514
	Management	TCP/541
	Contract Validation	TCP/443
FortiGate	HA Heartbeat	ETH Layer 0x8890, 0x8891, and 0x8893
	HA Synchronization	TCP/703, UDP/703
FortiGuard	AV/IPS Update	TCP/443, TCP/8890
	Cloud App DB	TCP/9582
	FortiGuard Queries	UDP/53, UDP/8888
	DNS	UDP/53, UDP/8888
	Registration	TCP/80
	Alert Email, Virus Sample	TCP/25
	Management, Firmware, SMS, FTM, Licensing, Policy Override	TCP/443
	Central Management, Analysis	TCP/541
FortiManager	Management	TCP/541
	IPv6	TCP/542
	Log & Report	TCP or UDP/514
	Secure SNMP	UDP/161, UDP/162
	FortiGuard Queries	TCP/8890, UDP/53
FortiSandbox	OFTP	TCP/514

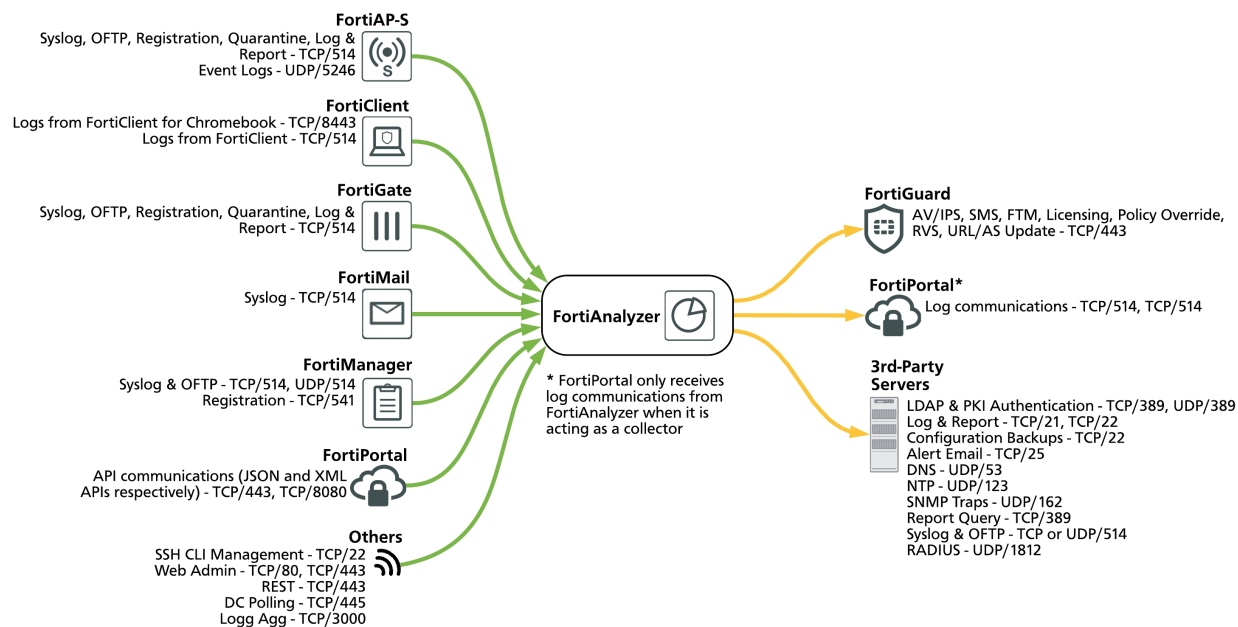
Outgoing Ports		
Purpose		Protocol/Port
Others	FSSO	TCP/8001 (by default; this port can be customized)



Note that, while a proxy is configured, FortiGate uses the following URLs to access the FortiGuard Distribution Network (FDN):

- **update.fortiguard.net**
- **service.fortiguard.net**
- **support.fortinet.com**

FortiAnalyzer open ports

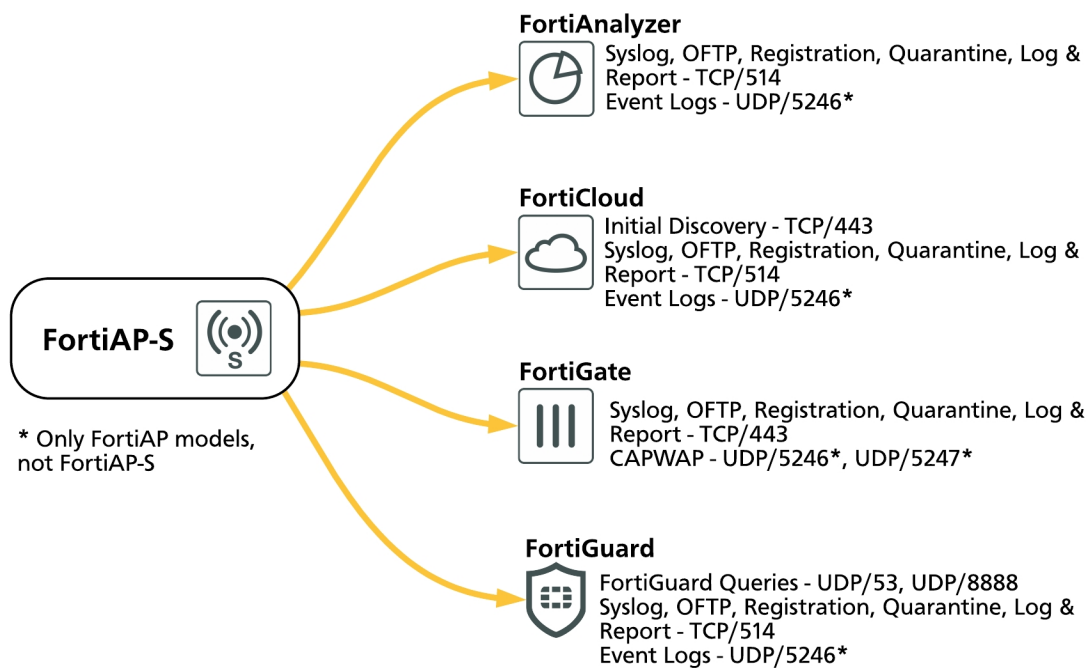


Incoming Ports		
Purpose		Protocol/Port
FortiAP-S	Syslog, OFTP, Registration, Quarantine, Log & Report	TCP/514
	Event Logs	UDP/5246
FortiClient	Logs from FortiClient for Chromebook	TCP/8443
	Logs from FortiClient (FortiClient must connect to FortiGate or EMS to send logs to FortiAnalyzer)	TCP/514
FortiGate	Syslog, OFTP, Registration, Quarantine, Log & Reports	TCP/514
FortiMail	Syslog	UDP/514
FortiManager	Syslog & OFTP	TCP/514, UDP/514
	Registration	TCP/541

Incoming Ports		
Purpose		Protocol/Port
FortiPortal	API communications (JSON and XML APIs respectively)	TCP/443, TCP/8080
Others	SSH CLI Management	TCP/22
	Web Admin	TCP/80, TCP/443
	REST	TCP/443
	DC Polling	TCP/445
	Logg Agg	TCP/3000

Outgoing Ports table		
Purpose		Protocol/Port
FortiGuard	AV/IPS, SMS, FTM, Licensing, Policy Override, RVS, URL/AS Update	TCP/443
FortiPortal (FortiPortal only receives log communications from FortiAnalyzer when it is acting as a collector)	Log communications	TCP/514, UDP/514
3rd-Party Servers	LDAP & PKI Authentication	TCP/389, UDP/389
	Log & Report	TCP/21, TCP/22
	Configuration Backups	TCP/22
	Alert Email	TCP/25
	DNS	UDP/53
	NTP	UDP/123
	SNMP Traps	UDP/162
	Report Query	TCP/389
	Syslog & OFTP	TCP or UDP/514
	RADIUS	UDP/1812

FortiAP-S open ports

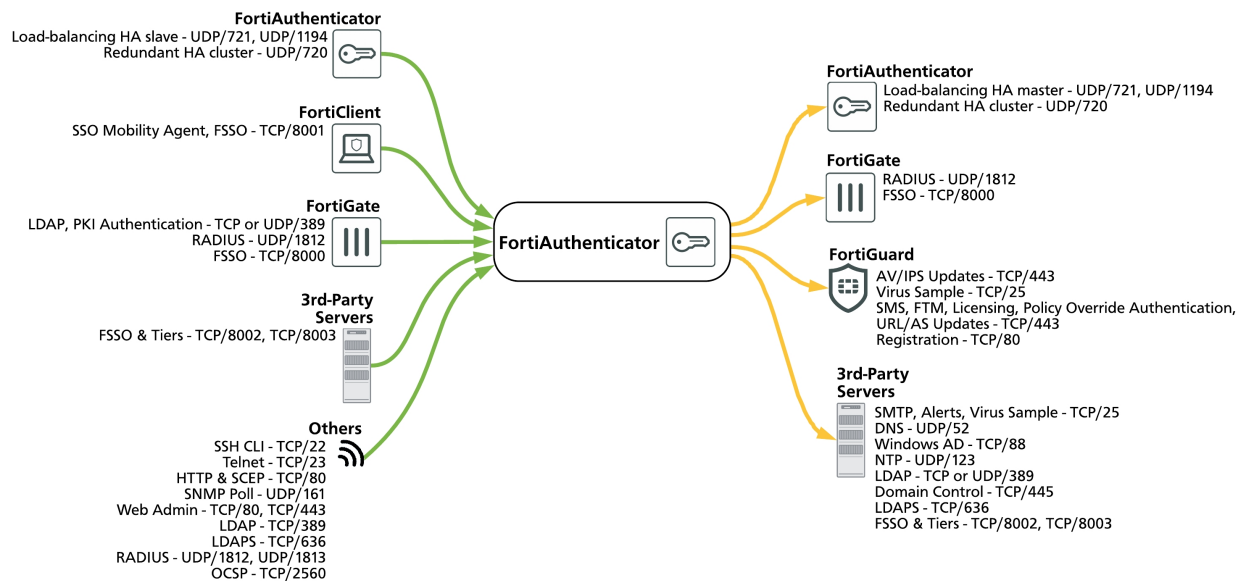


Outgoing Ports		
Purpose		Protocol/Port
FortiAnalyzer	Syslog, OFTP, Registration, Quarantine, Log & Report	TCP/514
	Event Logs	UDP/5246*
FortiCloud	Initial Discovery	TCP/443
	Syslog, OFTP, Registration, Quarantine, Log & Report	TCP/514
	Event Logs	UDP/5246*
FortiGate	Syslog, Registration, Quarantine, Log & Report	TCP/443
	CAPWAP	UDP/5246*, UDP/5247*

Outgoing Ports		
Purpose		Protocol/Port
FortiGuard	FortiGuard Queries	UDP/53, UDP/8888
	Syslog, OFTP, Registration, Quarantine, Log & Report	TCP/514
	Event Logs	UDP/5246*

* - Only FortiAP models, not FortiAP-S.

FortiAuthenticator open ports



Incoming Ports

Purpose	Protocol/Port
FortiAuthenticator	(HA) HA heartbeat
	LB slave sync
FortiClient	SSO Mobility Agent, FSSO
FortiGate	LDAP, PKI Authentication
	RADIUS
	FSSO

Incoming Ports		
Purpose		Protocol/Port
Others	SSH CLI	TCP/22
	Telnet	TCP/23
	HTTP & SCEP	TCP/80
	SNMP Poll	UDP/161
	Web Admin	TCP/80, TCP/443
	LDAP	TCP/389
	LDAPS	TCP/636
	RADIUS	UDP/1812, UDP/1813
	OCSP	TCP/2560
3rd-Party Servers	FSSO & Tiers	TCP/8002, TCP/8003

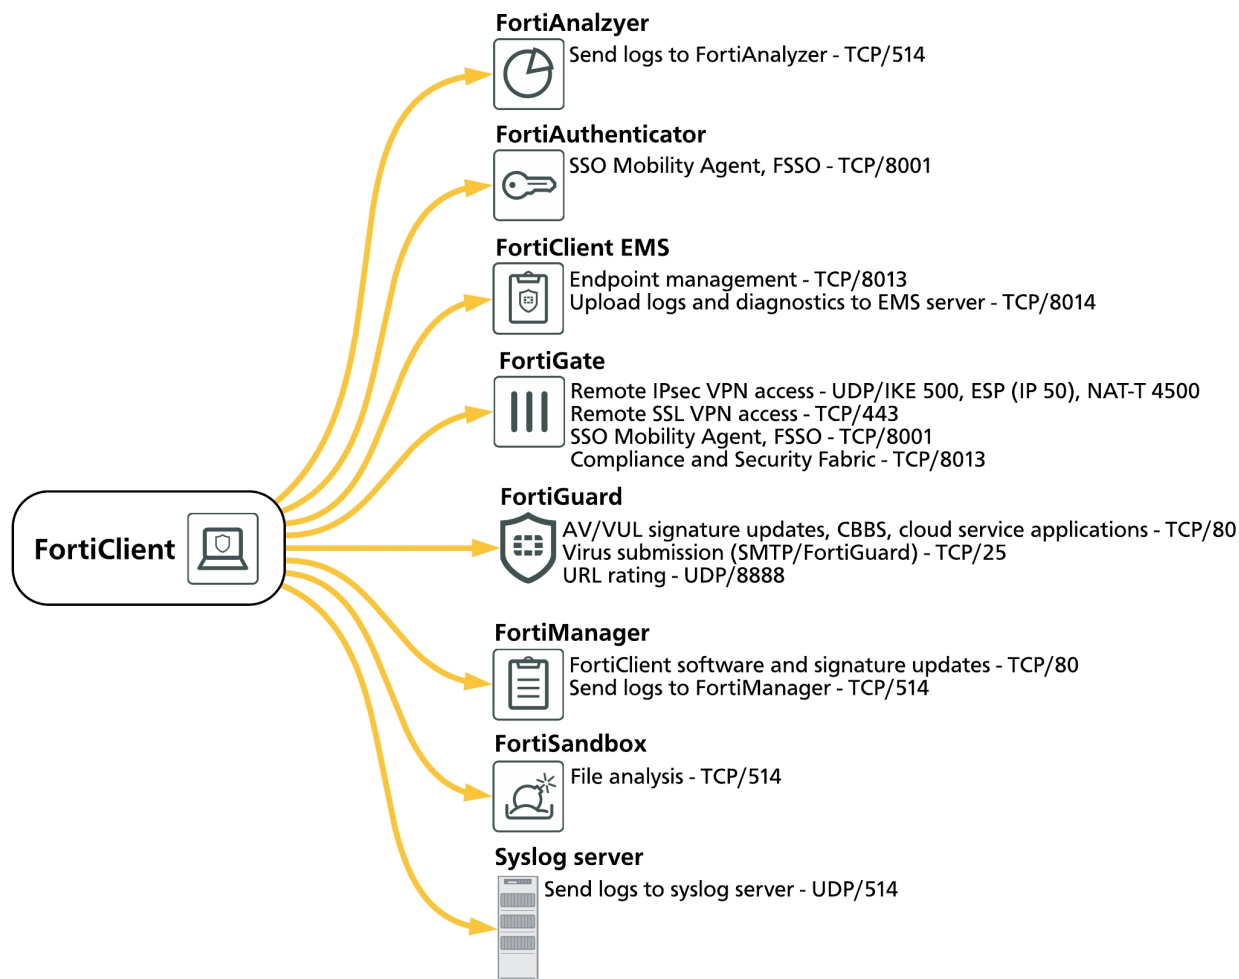
Outgoing Ports		
Purpose		Protocol/Port
FortiAuthenticator	(HA) HA heartbeat	UDP/720
	(LB slave) LB slave sync	UDP/721, UDP/1194
FortiGate	RADIUS	UDP/1812
	FSSO	TCP/8000
FortiGuard	AV/IPS Updates	TCP/443
	Virus Sample	TCP/25
	SMS, FTM, Licensing, Policy Override Authentication, URL/AS Updates	TCP/443
	Registration	TCP/80

Outgoing Ports		
Purpose		Protocol/Port
3rd-Party Servers	SMTP, Alerts, Virus Sample	TCP/25
	DNS	UDP/52
	Windows AD	TCP/88
	NTP	UDP/123
	LDAP	TCP or UDP389
	Domain Control	TCP/445
	LDAPS	TCP/636
	FSSO & Tiers	TCP/8002, TCP/8003

FortiClient open ports

The following diagrams and tables show the distinct communications for each FortiClient product.

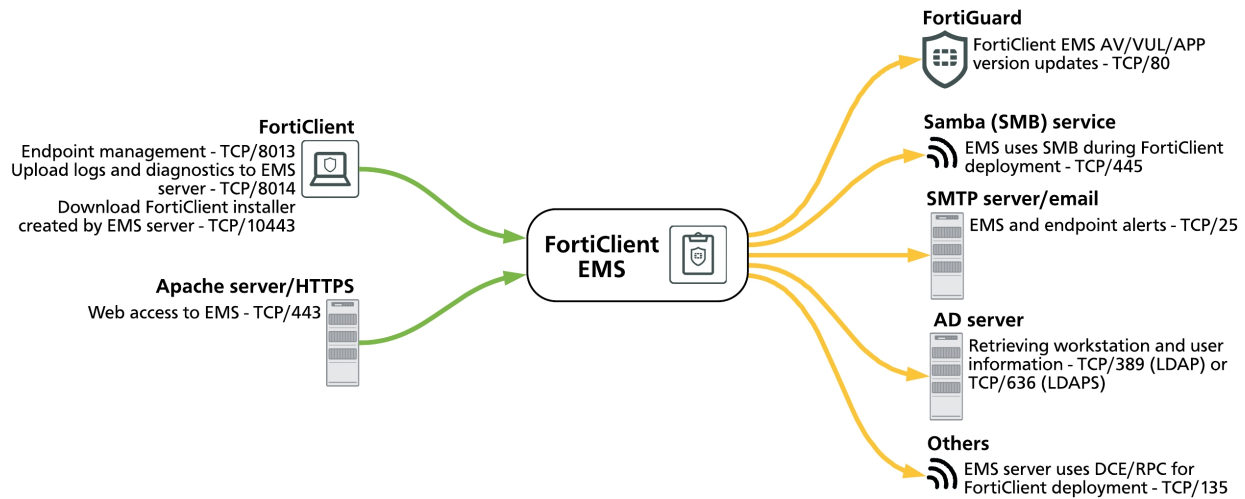
FortiClient



Outgoing Ports		
Purpose		Protocol/Port
FortiAnalyzer	Send logs to FortiAnalyzer (FortiClient must connect to FortiGate or EMS to send logs to FortiAnalyzer)	TCP/514
FortiAuthenticator	SSO Mobility Agent, FSSO	TCP/8001

Outgoing Ports		
Purpose		Protocol/Port
FortiClient EMS	Endpoint management	TCP/8013
	Upload logs and diagnostics to EMS server	TCP/8014
FortiGate	Remote IPsec VPN access	UDP/IKE 500, ESP (IP 50), NAT-T 4500
	Remote SSL VPN access	TCP/443 (by default; this port can be customized)
	SSO Mobility Agent, FSSO	TCP/8001
	Compliance and Security Fabric	TCP/8013 (by default; this port can be customized)
FortiGuard	AV/VUL signatures update, Cloud-based behavior scan (CBBS)/applications that use cloud services	TCP/80
	Virus submission (SMTP/FortiGuard)	TCP/25
	URL rating	UDP/8888 (by default; this port can be changed to port 53 via XML config file)
FortiManager	Select a FortiManager to be used for FortiClient signature updates	TCP/80 (by default; this port can be customized)
	Send logs to FortiManager (FortiClient must connect to FortiGate or EMS to send logs to FortiManager)	TCP/514
FortiSandbox	File analysis	TCP/514
Syslog server	Send logs to syslog server	UDP/514

FortiClient EMS



Incoming Ports

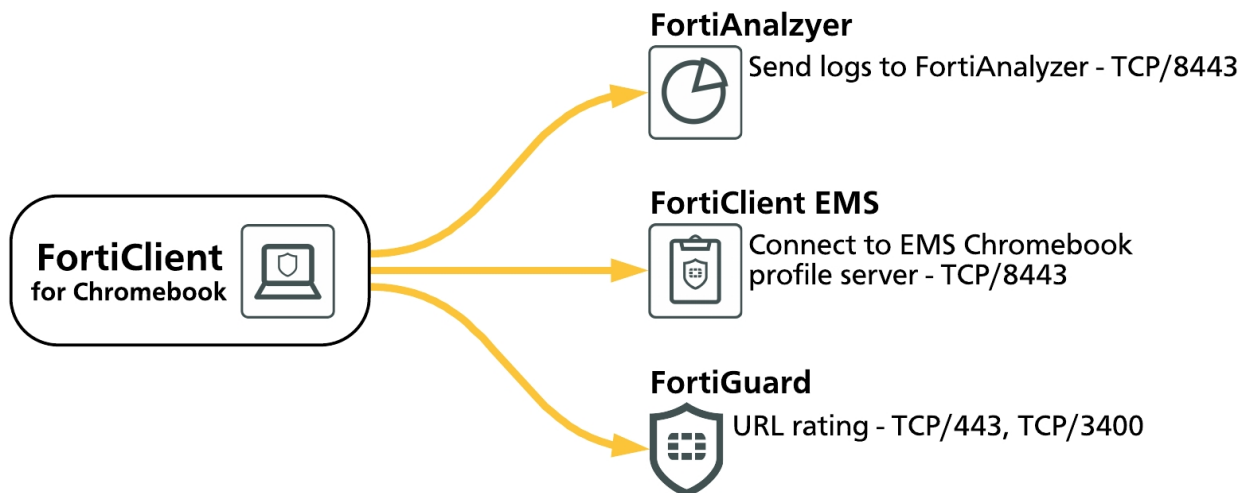
Purpose	Protocol/Port
FortiClient Endpoint management	TCP/8013 (by default; this port can be customized)
Upload logs and diagnostics to EMS server	TCP/8014
Download FortiClient installer created by EMS server	TCP/10443
Apache server/HTTPS Web access to EMS	TCP/443

Outgoing Ports

Purpose	Protocol/Port
FortiGuard FortiClient EMS AV/VUL/APP version updates	TCP/80
Samba (SMB) service EMS uses SMB during FortiClient deployment	TCP/445
SMTP server/email EMS and endpoint alerts	TCP/25

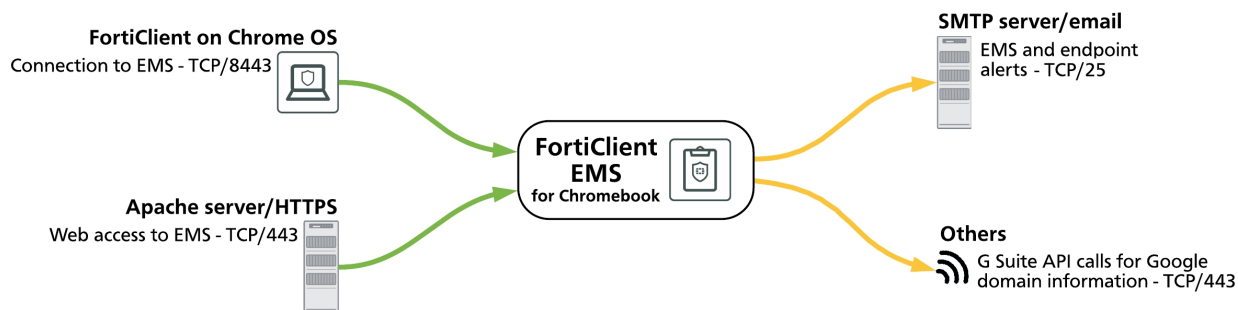
Outgoing Ports		
Purpose		Protocol/Port
AD server	Retrieving workstation and user information	TCP/389 or TCP/636 (for LDAP or LDAPS respectively)
Others	EMS server uses Distributed Computing Environment/Remote Procedure Calls (DCE/RPC) for FortiClient deployment	TCP/135

FortiClient for Chromebook



Outgoing Ports		
Purpose		Protocol/Port
FortiAnalyzer	Send logs to FortiAnalyzer	TCP/8443
FortiClient EMS	Connect to EMS Chromebook profile server	TCP/8443
FortiGuard	URL rating	TCP/443, TCP/3400

FortiClient EMS for Chromebook



Incoming Ports

Purpose	Protocol/Port
---------	---------------

FortiClient for Chromebook	Connection to EMS	TCP/8443
-----------------------------------	-------------------	----------

Apache server/HTTPS	Web access to EMS	TCP/443
----------------------------	-------------------	---------


Outgoing Ports

Purpose	Protocol/Port
---------	---------------

SMTP server/email	EMS and endpoint alerts	TCP/25
--------------------------	-------------------------	--------

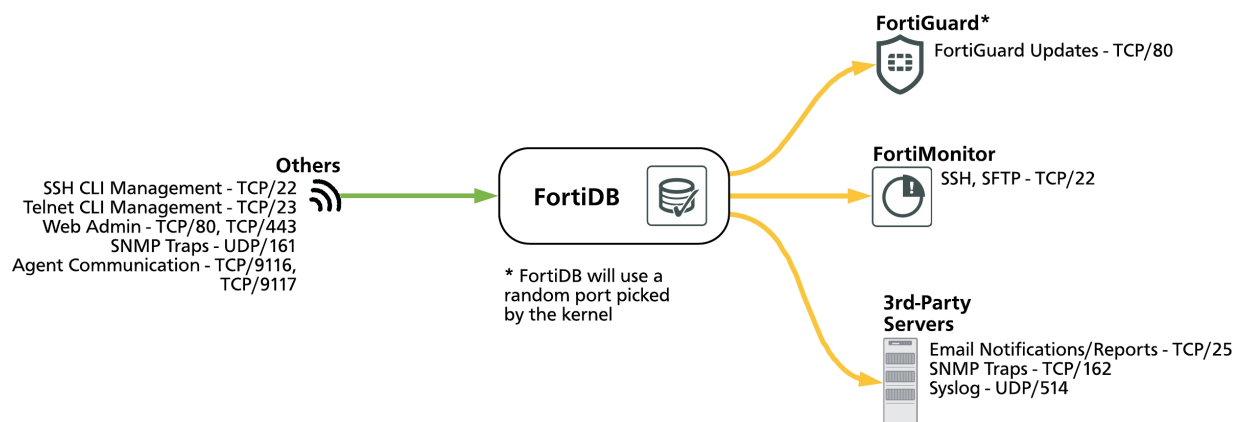
Others	G Suite API calls for Google domain information	TCP/443
---------------	---	---------

FortiCloud open ports



Incoming Ports		
Purpose		Protocol/Port
FortiAP-S	Initial Discovery	TCP/443
	Syslog, OFTP, Registration, Quarantine, Log & Report	TCP/514
	Event Logs	UDP/5246
FortiGate	Registration, Quarantine, Log & Report, Syslog	TCP/443
	OFTP	TCP/514
	Management	TCP/541
	Contract Validation	TCP/443
Outgoing Ports		
Purpose		Protocol/Port
FortiGuard	Registration	TCP/443

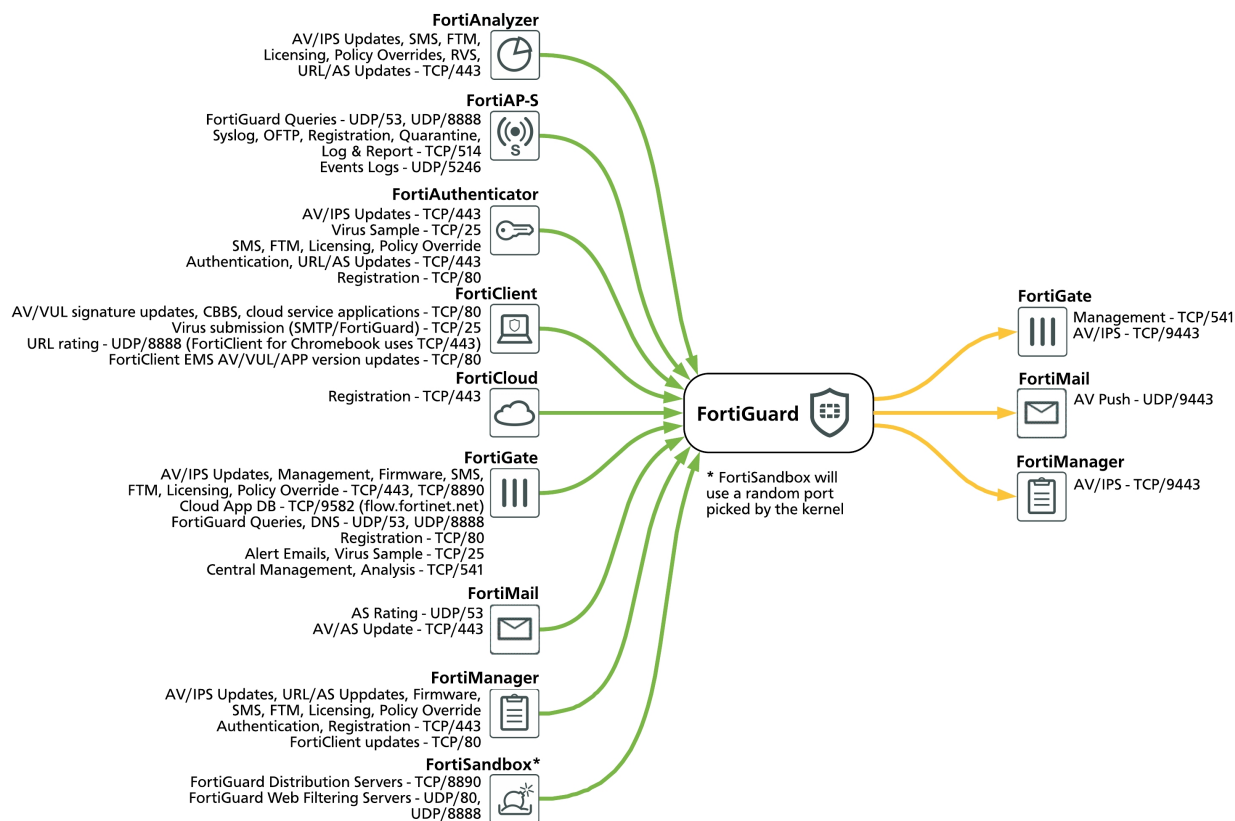
FortiDB open ports



Incoming Ports		
Purpose		Protocol/Port
Others	SSH CLI Management	TCP/22
	Telnet CLI Management	TCP/23
	Web Admin	TCP/80, TCP/443
	SNMP Traps	UDP/161
	Agent Communication	TCP/9116, TCP/9117

Outgoing Ports		
Purpose		Protocol/Port
FortiGuard (FortiDB will use a random port picked by the kernel)	FortiGuard Updates	TCP/80
FortiMonitor	SSH, SFTP	TCP/22
3rd-Party Servers	Email Notifications/Reports	TCP/25
	SNMP Traps	UDP/162
	Syslog	UDP/514

FortiGuard open ports



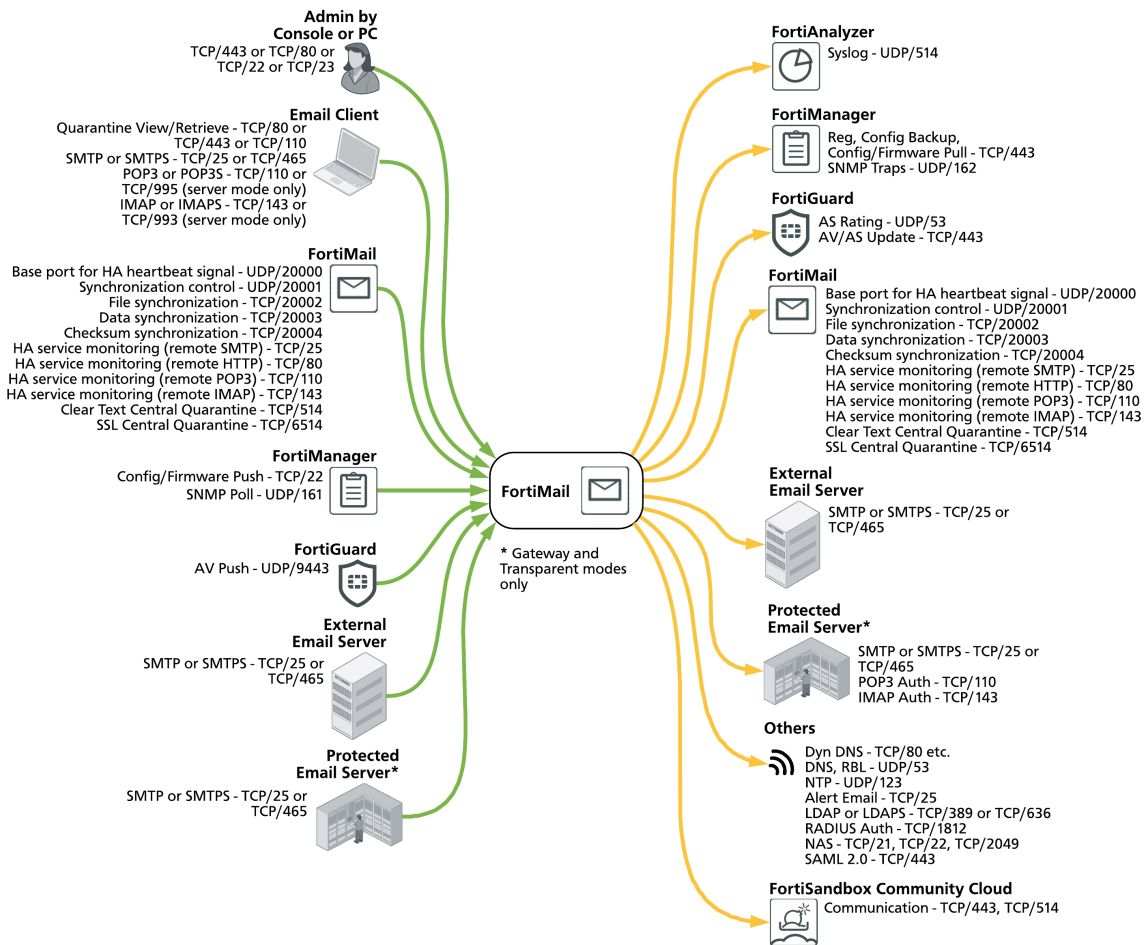
Incoming Ports

Purpose	Protocol/Port
FortiAnalyzer	AV/IPS Updates, SMS, FTM, Licensing, Policy Overrides, RVS, URL/AS Update TCP/443
FortiAP-S	FortiGuard Queries UDP/53, UDP/8888
	Syslog, OFTP, Registration, Quarantine, Log & Report TCP/514
	Event Logs UDP/5246

Incoming Ports		
Purpose		Protocol/Port
FortiAuthenticator	AV/IPS Updates	TCP/443
	Virus Sample	TCP/25
	SMS, FTM, Licensing, Policy Override Authentication, URL/AS Updates	TCP/443
	Registration	TCP/80
FortiClient	AV/VUL signatures update, Cloud-based behavior scan (CBBS)/applications that use cloud services	TCP/80
	Virus submission (SMTP/FortiGuard)	TCP/25
	URL rating	UDP/8888 (by default; this port can be changed to port 53 via XML config file) Note: FortiClient for Chromebooks contacts FortiGuard for URL ratings via TCP/443
	FortiClient EMS AV/VUL/APP version updates *	TCP/80
FortiCloud	Registration	TCP/443
FortiGate	AV/IPS Update, Management, Firmware, SMS, FTM, Licensing, Policy Override	TCP/443, TCP/8890
	Cloud App DB	TCP/9582 (flow.fortinet.net)
	FortiGuard Queries, DNS	UDP/53, UDP/8888
	Registration	TCP/80
	Alert Emails, Virus Sample	TCP/25
	Central Management, Analysis	TCP/541
FortiMail	AS Rating	UDP/53
	AV/AS Update	TCP/443

Incoming Ports		
Purpose		Protocol/Port
FortiManager	AV/IPS Updates, URL/AS Update, Firmware, SMS, FTM, Licensing, Policy Override Authentication, Registration	TCP/443
	FortiClient updates	TCP/80
FortiSandbox (FortiSandbox will use a random port picked by the kernel)	FortiGuard Distribution Servers	TCP/8890
	FortiGuard Web Filtering Servers	UDP/53, UDP/8888
Outgoing Ports		
Purpose		Protocol/Port
FortiGate	Management	TCP/541
	AV/IPS	UDP/9443
FortiMail	AV Push	UDP/9443
FortiManager	AV/IPS	UDP/9443

FortiMail open ports



When operating in its default configuration, FortiMail does not accept TCP or UDP connections on any port except port1 and port2 network interfaces, which accept:

- ICMP pings,
- HTTPS connections on TCP/443,
- and SSH connections on TCP/22.

Incoming Ports		
Purpose		Protocol/Port
Admin by Console or PC	SSH, Telnet, HTTP, SSH, Console	TCP/443 or TCP/80 or TCP/22 or TCP/23
Email Client	Quarantine View/Retrieve	TCP/80 or TCP/443 or TCP/110
	SMTP or SMTPS	TCP/25 or TCP/465
	POP3 or POP3S	TCP/110 or TCP/995 (server mode only)
	IMAP or IMAPS	TCP/143 or TCP/993 (server mode only)
	WebDAV and CalDAV	TCP/8008
FortiMail	Base port for HA heartbeat signal	UDP/20000
	Synchronization control	UDP/20001
	File synchronization	TCP/20002
	Data synchronization	TCP/20003
	Checksum synchronization	TCP/20004
	HA service monitoring (remote SMTP)	TCP/25
	HA service monitoring (remote HTTP)	TCP/80
	HA service monitoring (remote POP3)	TCP/110
	HA service monitoring (remote IMAP)	TCP/143
	Clear Text Central Quarantine	TCP/514
	SSL Central Quarantine	TCP/6514
FortiManager	SNMP Poll	TCP/161
	AV Push	
FortiGuard	AV Push	UDP/9443

Incoming Ports		
Purpose		Protocol/Port
External Email Server	SMTP or SMTPS	TCP/25 or 465
	Storage: iSCSI, NFS	TCP/3260 (iSCSI), TCP/2049 (NFS)
	Config Backup	SFTP / FTP
	Mail Data Backup	NFS, SMB/CIFS, SSH, external USB (direct connected), iSCSI
Protected Email Server	SMTP or SMTPS	TCP/25 or 465

Outgoing Ports		
Purpose		Protocol/Port
FortiAnalyzer	OFTP	UDP/514
FortiManager	SNMP Traps	UDP/162
	AV/AS Query	
FortiGuard	AS Rating	UDP/53 or 8888, 8889
	AV/AS Update	TCP/443

Outgoing Ports		
Purpose		Protocol/Port
FortiMail	Base port for HA heartbeat signal	UDP/20000
	Synchronization control	UDP/20001
	File synchronization	TCP/20002
	Data synchronization	TCP/20003
	Checksum synchronization	TCP/20004
	HA service monitoring (remote SMTP)	TCP/25
	HA service monitoring (remote HTTP)	TCP/80
	HA service monitoring (remote POP3)	TCP/110
	HA service monitoring (remote IMAP)	TCP/143
	Clear Text Central Quarantine	TCP/514
	SSL Central Quarantine	TCP/6514
External Email Server	SMTP or SMTPS	TCP/25 or TCP/465
Protected Email Server	SMTP or SMTPS	TCP/25 or TCP/465
	POP3 Auth	TCP/110
	IMAP Auth	TCP/143
Others	Dyn DNS	TCP/80 *
	DNS, RBL	UDP/53
	NTP	UDP/123
	Alert Email	TCP/25
	LDAP or LDAPS	TCP/389 or TCP/636
	RADIUS Auth	TCP/1812
	NAS	TCP/21, TCP/22, TCP/2049
	OCSP (for PKI user)	TCP/80, or defined by certificate

Outgoing Ports		
Purpose		Protocol/Port
FortiSandbox / FortiSandbox Cloud	Communication	TCP/443, TCP/514

* - FortiMail generates outbound traffic and sends an HTTP SYN request via TCP/80. The Fortinet RSS Feed widget provides a convenient display of the latest security advisories and discovered threats from Fortinet. Also, if an email message contains a shortened URI that redirects to another URI, it would cause FortiMail to send an HTTP SYN request to the shortened URI to get the redirected URI.



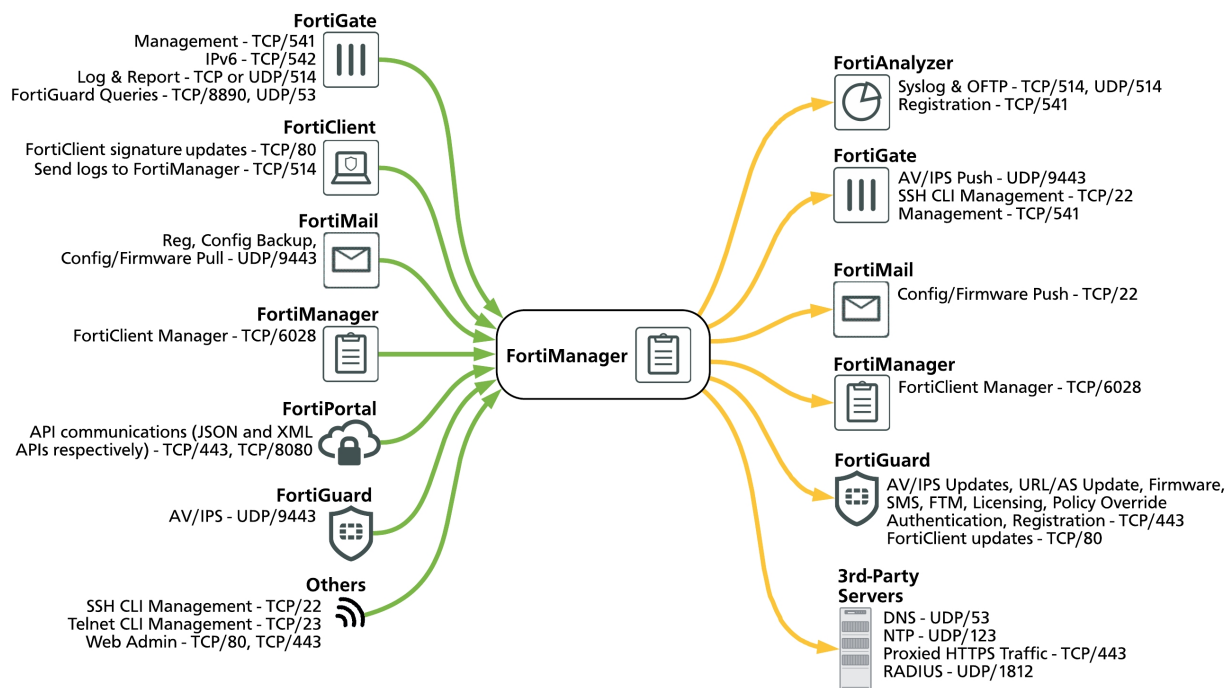
Note that FortiMail uses the following URLs to access the FortiGuard Distribution Network (FDN):

- **update.fortiguard.net**
- **service.fortiguard.net**
- **support.fortinet.com**

Furthermore, FortiMail performs these queries and updates listed below using the following ports and protocols:

- FortiGuard Anti-Spam rating queries: UDP/53, 8888, 8889
- FortiGuard AntiVirus Push updates: UDP/9443
- FortiGuard Anti-Spam or AntiVirus updates: TCP/443

FortiManager open ports



Incoming Ports		
Purpose		Protocol/Port
FortiGate	Management	TCP/541
	IPv6	TCP/542
	Log & Report	TCP or UDP/514
	FortiGuard Queries	TCP/8890, UDP/53
FortiClient	Select a FortiManager to be used for FortiClient signature updates	TCP/80 (by default; this port can be customized)
	Send logs to FortiManager (FortiClient must connect to FortiGate or EMS to send logs to FortiManager)	TCP/514
FortiGuard	AV/IPS	UDP/9443

Incoming Ports		
Purpose		Protocol/Port
FortiMail	Registration	UDP/9443
	AV/AS Query	
FortiManager	FortiClient Manager	TCP/6028
FortiPortal	API communications (JSON and XML APIs respectively)	TCP/443, TCP/8080
Others	SSH CLI Management	TCP/22
	Telnet CLI Management	TCP/23
	Web Admin	TCP/80, TCP/443

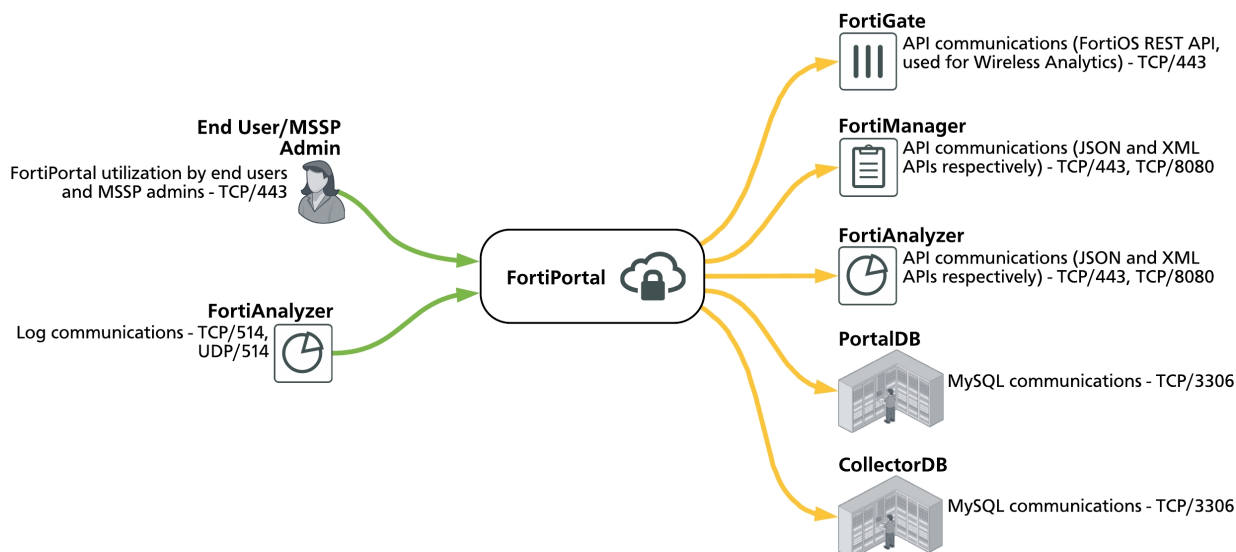
Outgoing Ports		
Purpose		Protocol/Port
FortiAnalyzer	Syslog & OFTP	TCP/514, UDP/514
	Registration	TCP/541
FortiGate	AV/IPS Push	UDP/9443
	SSH CLI Management	TCP/22
	Management	TCP/541
FortiGuard	AV/IPS Updates, URL/AS Update, Firmware, SMS, FTM, Licensing, Policy Override Authentication, Registration	TCP/443
	FortiClient updates	TCP/80
FortiMail	AV Push	
FortiManager	FortiClient Manager	TCP/6028
3rd-Party Servers	DNS	UDP/53
	NTP	UDP/123
	Proxied HTTPS Traffic	TCP/443
	RADIUS	UDP/1812



Note that, while a proxy is configured, FortiManager uses the following URLs to access the FortiGuard Distribution Network (FDN) for the following updates:

- **fds1.fortinet.com** - FortiGate AV/IPS package downloads
- **guard.fortinet.net** - Webfilter/Anti-Spam DB and AVfileQuery DB downloads
- **forticlient.fortinet.com** - FortiClient signature package downloads
- **fgd1.fortigate.com:8888** - FortiClient Webfilter queries to FortiGuard

FortiPortal open ports



Incoming Ports

Purpose

Protocol/Port

End User/MSSP Admin

FortiPortal utilization by end users and MSSP admins

TCP/443

FortiAnalyzer
(FortiPortal only receives log communications from FortiAnalyzer when it is acting as a collector)

Log communications

TCP/514, UDP/514

Outgoing Ports table

Purpose

Protocol/Port

FortiGate

API communications (FortiOS REST API, used for Wireless Analytics)

TCP/443

FortiManager

API communications (JSON and XML APIs respectively)

TCP/443, TCP/8080


FortiAnalyzer

API communications (JSON and XML APIs respectively)

TCP/443, TCP/8080

Outgoing Ports table		
Purpose		Protocol/Port
PortalDB	MySQL communications	TCP/3306
CollectorDB	MySQL communications	TCP/3306

FortiSandbox open ports



Incoming Ports		
Purpose		Protocol/Port
FortiGate	OFTP	TCP/514
FortiClient	File analysis	TCP/514
Others	SSH CLI Management	TCP/22
	Telnet CLI Management	TCP/23
	Web Admin	TCP/80, TCP/443
	OFTP Communication with FortiGate & FortiMail	TCP/514
	Third-party proxy server for ICAP servers	ICAP: TCP/1344 ICAPS: TCP/11344

Outgoing Ports		
Purpose		Protocol/Port
FortiGuard (FortiSandbox will use a random port picked by the kernel)	FortiGuard Distribution Servers	TCP/8890
	FortiGuard Web Filtering Servers	UDP/53, UDP/8888

Outgoing Ports		
Purpose		Protocol/Port
FortiSandbox Community Cloud (FortiSandbox will use a random port picked by the kernel)	Upload detected malware information	TCP/443, UDP/53



Note that FortiSandbox uses the following FQDNs to access the FortiSandbox Community Cloud, depending on which protocol and port is used:

- TCP/443: **fqdl.fortinet.net**
- UDP/53: **fqsrv.fortinet.net**

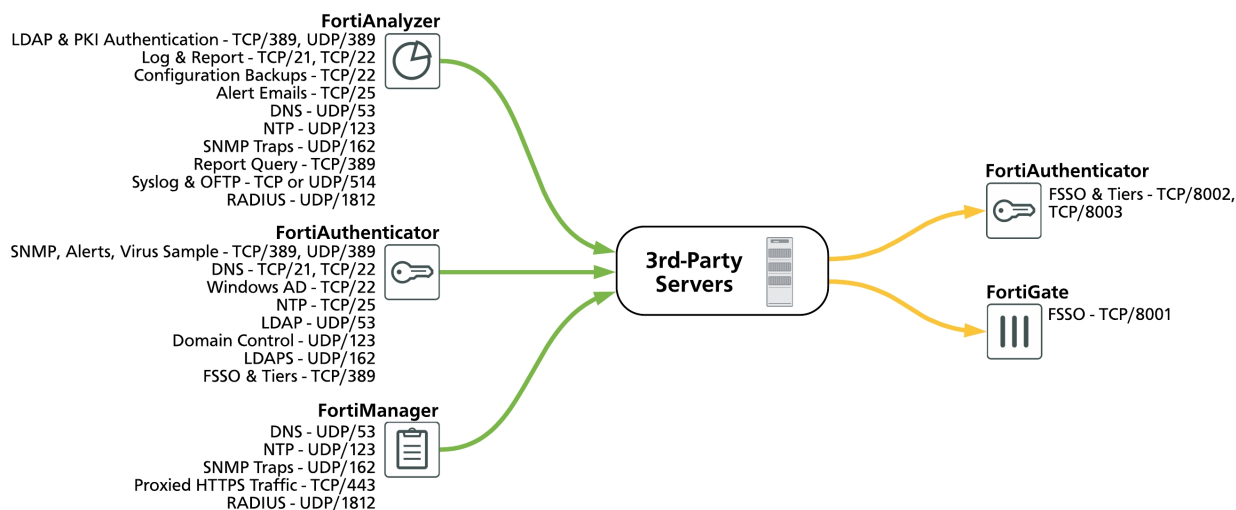
Services and port numbers required for FortiSandbox

The tables above show all the services required for FortiSandbox to function correctly. You can use the diagnostic FortiSandbox command `test-network` to verify that all the services are allowed by the upstream. If the result is `Passed`, then there is no issue. If there is an issue with a specific service, it will be shown in the command output, and inform you which port needs to be opened.

This command checks:

- VM Internet access
- Internet connection
- System DNS resolve speed
- VM DNS resolve speed
- Ping speed
- Wget speed
- Web Filtering service
- FortiSandbox Community Cloud service

3rd-party servers open ports



Incoming Ports		
Purpose		Protocol/Port
FortiAnalyzer	LDAP & PKI Authentication	TCP/389, UDP/389
	Log & Report	TCP/21, TCP/22
	Configuration Backups	TCP/22
	Alert Emails	TCP/25
	DNS	UDP/53
	NTP	UDP/123
	SNMP Traps	UDP/162
	Report Query	TCP/389
	Syslog & OFTP	TCP or UDP/514
	RADIUS	UDP/1812

Incoming Ports		
Purpose		Protocol/Port
FortiAuthenticator	SMTP, Alerts, Virus Sample	TCP/25
	DNS	UDP/52
	Windows AD	TCP/88
	NTP	UDP/123
	LDAP	TCP or UDP/389
	Domain Control	TCP/445
	LDAPS	TCP/636
	FSSO & Tiers	TCP/8002, TCP/8003
FortiManager	DNS	UDP/53
	NTP	UDP/123
	SNMP Traps	UDP/162
	Proxied HTTPS Traffic	TCP/443
	RADIUS	UDP/1812

Outgoing Ports		
Purpose		Protocol/Port
FortiAuthenticator	FSSO & Tiers	TCP/8002, TCP/8003
FortiGate	FSSO	TCP/8001 (by default; this port can be customized)

Fortinet proprietary protocols

The following section provides a full list of Fortinet's proprietary protocols, their purposes, and what ports they operate on:

- FGCP - FortiGate Clustering Protocol
- FGSP - FortiGate Session Life Support Protocol
- FGFM - FortiGate to FortiManager Protocol
- SLBC - Session-aware Load Balancing Cluster
- Fortinet Security Fabric
- FortiTelemetry/On-Net/FortiClient Endpoint Compliance
- FortiGuard
- FortiLink
- FortiOS WAN optimization
- FSSO - Fortinet Single Sign-On
- OFTP - Optimized Fabric Transfer Protocol
- FortiClient EMS - Enterprise Management Server

FGCP - FortiGate Clustering Protocol

In an active-passive HA configuration, the FortiGate Clustering Protocol (FGCP) provides failover protection, whereby the cluster can provide FortiGate services even when one of the cluster units loses connection. FGCP is also a Layer 2 heartbeat that specifies how FortiGate units communicate in an HA cluster and keeps the cluster operating.



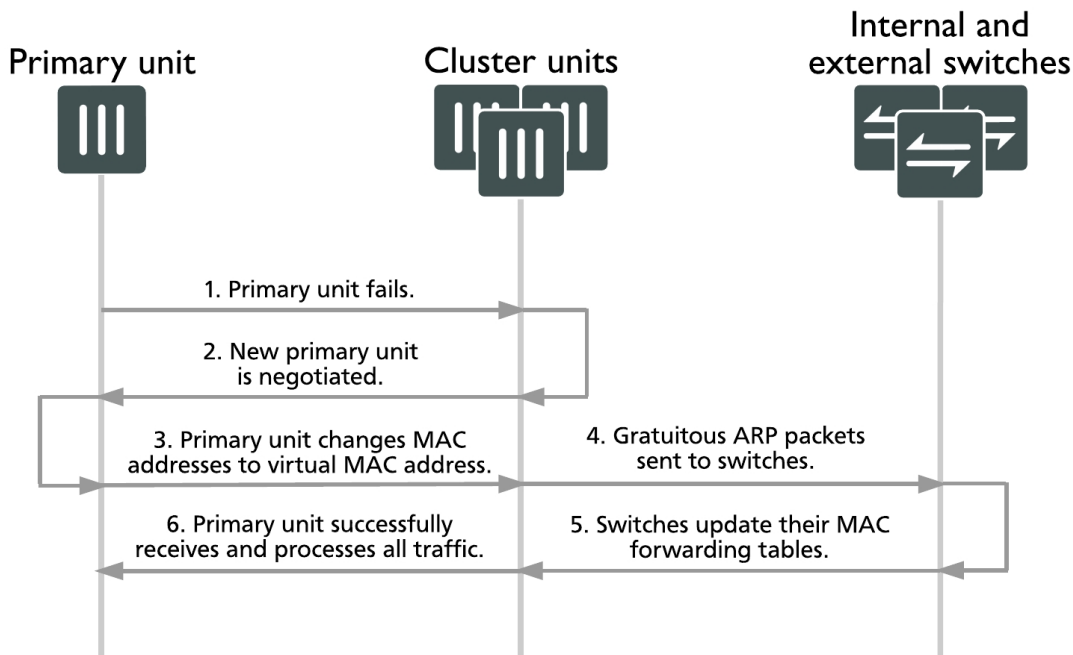
You cannot mix FGCP and SLBC clusters in the same chassis.

The FortiGate's HA Heartbeat listens on ports TCP/703, TCP/23, or ETH Layer 2/8890.

Virtual MAC addresses

FGCP assigns virtual MAC addresses to each primary unit interface in an HA cluster. Virtual MAC addresses are in place so that, if a failover occurs, the new primary unit interfaces will have the same MAC addresses as the failed primary unit interfaces. If the MAC addresses were to change after a failover, the network would take longer to recover because all attached network devices would have to learn the new MAC addresses before they could communicate with the cluster.

If a cluster is operating in Transparent mode, FGCP assigns a virtual MAC address for the primary unit management IP address. Since you can connect to the management IP address from any interface, all of the FortiGate interfaces appear to have the same virtual MAC address.



When a cluster starts up, after a failover, the primary unit sends gratuitous ARP packets to update the switches connected to the cluster interfaces with the virtual MAC address. The switches update their MAC forwarding tables with this MAC address. As a result, the switches direct all network traffic to the primary unit. Depending on the cluster configuration, the primary unit either processes this network traffic itself or load balances the network traffic among all of the cluster units.

You cannot disable sending gratuitous ARP packets, but you can change the number of packets that are sent (1-60 ARP packets) by entering the following command:

```
config system ha
    set arps <integer>
end
```

You can change the time between ARP packets (1-20 seconds) by entering the following command:

```
config system ha
    set arps-interval <integer>
end
```

Assigning virtual MAC addresses

Virtual MAC addresses are determined based on the following formula:

00-09-0f-09-<group-id_hex>-<vcluster_integer><idx>

where:

- **<group-id_hex>**: The HA group ID for the cluster converted to hexadecimal. The table below lists some example virtual MAC addresses set for each group ID:

Integer Group ID	Hexadecimal Group ID
0	00
1	01
2	02
3	03
...	...
10	0a
11	0b
...	...
63	3f
...	...
255	ff

- **<vcluster_integer>**: This value is 0 for virtual cluster 1 and 2 for virtual cluster 2. If virtual domains are not enabled, HA sets the virtual cluster to 1 and by default all interfaces are in the root virtual domain. Including virtual cluster and virtual domain factors in the virtual MAC address formula means that the same formula can be used whether or not virtual domains and virtual clustering is enabled.
- **<idx>**: The index number of the interface. In NAT/Route mode, interfaces are numbered from 0 to x (where x is the number of interfaces). The interfaces are listed in alphabetical order on the web-based manager and CLI. The interface at the top of the interface list is first in alphabetical order by name and has an index of 0. The second interface in the list has an index of 1 and so on. In Transparent mode, the index number for the management IP address is 0.

Every FortiGate unit physical interface has two MAC addresses: the current hardware address and the permanent hardware address. The permanent hardware address cannot be changed, as it is the actual MAC address of the interface hardware. The current hardware address can be changed, but only when a FortiGate unit is **not** operating in HA. For an operating cluster, the current hardware address of each cluster unit interface is changed to the HA virtual MAC address by the FGCP.

You cannot change an interface MAC address and you cannot view MAC addresses from the system interface CLI command.

You can use the `get hardware nic <interface_name_str>` (or `diagnose hardware deviceinfo nic <interface_str>`) command to display both MAC addresses for any FortiGate interface. This command displays hardware information for the specified interface, including the current hardware address (as `Current_HWaddr`) and the permanent hardware address (as `Permanent_HWaddr`). For some interfaces, the current hardware address is displayed as `MAC`.

Failover protection

FGCP supports three kinds of failover protection:

1. **Device failover:** Automatically replaces a failed device and restarts traffic flow with minimal impact on the network. All subordinate units in an active-passive HA cluster are constantly waiting to negotiate to become primary units. Only the heartbeat packets sent by the primary unit keep the subordinate units from becoming primary units. Each received heartbeat packet resets negotiation timers in the subordinate units. If this timer is allowed to run out because the subordinate units do not receive heartbeat packets from the primary unit, the subordinate units assume that the primary unit has failed, and negotiate to become primary units themselves. The default time interval between HA heartbeats is 200 ms.
2. **Link failover:** Maintains traffic flow if a link fails. In this case, the primary unit does not stop operating, and therefore participates in the negotiation of selecting a new primary unit. The old primary unit then joins the cluster as a subordinate unit. Furthermore, any subordinate units with a link failure are unlikely to become the primary unit in future negotiations.
3. **Session failover:** With session failover (also called session pickup) enabled, the primary unit informs the subordinate units of changes to the primary unit connection and state tables, keeping the subordinate units up-to-date with the traffic currently being processed by the cluster. This helps new primary units resume communication sessions with minimal loss of data, avoiding the need to restart active sessions.

Synchronization of configurations

The FGCP uses a combination of incremental and periodic synchronization to make sure that the configuration of all cluster units is synchronized to that of the primary unit. However, there are certain settings that are not synchronized between cluster units:

- HA override
- HA device priority


- The virtual cluster priority
- The FortiGate unit host name
- The HA priority setting for a ping server (or dead gateway detection) configuration
- The system interface settings of the HA reserved management interface
- The HA default route for the reserved management interface, set using the `ha-mgmt-interface-gateway` option of the `config system ha` command.

You can disable configuration synchronization by entering the following command:

```
config system ha
    set sync-config disable
end
```

The command `execute ha synchronize` can be used to perform a manual synchronization.

The FGCP heartbeat operates on TCP port 703 with an independent IP address not assigned to any FortiGate interface. You can create an FGCP cluster of up to four FortiGate units. Below is an example of FGCP used to create an HA cluster installed between an internal network and the Internet.



FGCP HA provides a solution for two key requirements of critical enterprise networking: enhanced reliability and increased performance, through device, link, and remote link failover protection. Extended FGCP features include full mesh HA and virtual clustering. You can also fine tune the performance of the FGCP to change how a cluster forms and shares information among cluster units and how the cluster responds to failures.

Before configuring an FGCP HA cluster, make sure your FortiGate interfaces are configured with static IP addresses. If any interface gets its address using DHCP or PPPoE you should temporarily switch it to a static address and enable DHCP or PPPoE after the cluster has been established.




Heartbeat traffic, such as FGCP, uses multicast on port number 6065 and uses link-local IPv4 addresses in the 169.254.0.x range. HA heartbeat packets have an Ethertype field value of **0x8890**.

Synchronization traffic, such as FGSP, uses unicast on port number 6066 and the IP address 239.0.0.2. HA sessions that synchronize the cluster have an Ethertype field value of **0x8893**.

The HA IP addresses are hard-coded and cannot be configured.

How to set up FGCP clustering

This example describes how to enhance the reliability of a network protected by a FortiGate unit by adding a second FortiGate unit to create a FortiGate Clustering Protocol (FGCP) HA cluster. The FortiGate already on the network will be configured to become the primary unit by increasing its device priority and enabling override. The new FortiGate will be prepared by setting it to factory defaults to wipe any configuration changes. Then it will be licensed, configured for HA, and then connected to the FortiGate already on the network. The new FortiGate becomes the backup unit and its configuration is overwritten by the primary unit.



If you have not already done so, register the primary FortiGate and apply licenses to it before setting up the cluster. This includes FortiCloud activation and FortiClient licensing, and entering a license key if you purchased more than 10 Virtual Domains (VDOMs). You can also install any third-party certificates on the primary FortiGate before forming the cluster.

The FortiGates should be running the same FortiOS firmware version, and their interfaces should not be configured to get their addresses from DHCP or PPPoE.

Configuring the primary FortiGate

1. Connect to the primary FortiGate and go to **Dashboard > Main > System Information**. Change the unit's **Host Name** to identify it as the primary FortiGate.

You can also enter this CLI command:

```
config system global
  set hostname Primary_FortiGate
end
```

2. You then need to set the HA mode to active-passive. Enter the following CLI command to set the HA mode to active-passive, set a group name and password, increase the device priority to a higher value (for example, 250) and enable override:

```
config system ha
  set mode a-p
  set group-name My-HA-Cluster
```



```
set password
set priority 250
set override enable
set hbdev ha1 50 ha2 50
end
```

This command also selects ha1 and ha2 to be the heartbeat interfaces, with their priorities set to 50. Enabling override and increasing the priority ensures that this FortiGate should become the primary unit.



You can configure these settings in the GUI under **System > HA**, however the override can *only* be enabled in the CLI.

Configuring the backup FortiGate

1. Enter the CLI command below to reset the new FortiGate to factory default settings (skip this step if the FortiGate is fresh from the factory). It is recommended to set it back to factory defaults to reduce the chance of synchronization problems.:

```
execute factoryreset
```

2. Make sure to change the firmware running on the new FortiGate to the same version running on the primary unit, register, and apply licenses to it before adding it to the cluster.
3. Then go to **Dashboard > Main > System Information**. Change the unit's **Host Name** to identify it as the backup FortiGate.

You can also enter this CLI command:

```
config system global
  set hostname Backup_FortiGate
end
```

4. Duplicate the primary unit's HA settings, except make sure to set the backup device's priority to a lower value and do *not* enable override.

Connecting the cluster

Connect the HA cluster as shown in the initial diagram above. Making these connections will disrupt network traffic as you disconnect and re-connect cables.

When connected, the primary and backup FortiGates find each other and negotiate to form an HA cluster. The primary unit synchronizes its configuration with the backup FortiGate. Forming the cluster happens automatically with minimal or no disruption to network traffic.

Heartbeat packet ethertypes

Normal IP packets are 802.3 packets that have an ethernet type (ethertype) field value of 0x0800. Ethertype values other than 0x0800 are understood as level 2 frames rather than IP packets.

By default, HA heartbeat packets use the following ethertypes:

- HA heartbeat packets for NAT/Route mode clusters use Ethertype 0x8890. These packets are used by cluster units to find other cluster units and to verify the status of other cluster units while the cluster is operating. You can change the Ethertype of these packets using the `ha-eth-type` option under `config system ha`.
- HA heartbeat packets for Transparent mode clusters use Ethertype 0x8891. These packets are used by cluster units to find other cluster units and to verify the status of other cluster units while the cluster is operating. You can change the Ethertype of these packets using the `hc-eth-type` option under `config system ha`.

- HA telnet sessions between cluster units over HA heartbeat links use Ethertype 0x8893. The telnet sessions allow an administrator to connect between FortiGates in the cluster using the `execute ha manage` command. You can change the Ethertype of these packets using the `l2ep-eth-type` option under `config system ha`.

Because heartbeat packets are recognized as level 2 frames, the switches and routers on your heartbeat network that connect to heartbeat interfaces must be configured to allow them. If level2 frames are dropped by these network devices, heartbeat traffic will not be allowed between the cluster units.

Some third-party network equipment may use packets with these Ethernets for other purposes. For example, Cisco N5K/Nexus switches use Ethertype 0x8890 for some functions. When one of these switches receives Ethertype 0x8890 packets from an attached cluster unit, the switch generates CRC errors and the packets are not forwarded. As a result, FortiGate units connected with these switches cannot form a cluster.

In some cases, if the heartbeat interfaces are connected and configured so regular traffic flows but heartbeat traffic is not forwarded, you can change the configuration of the switch that connects the HA heartbeat interfaces to allow level2 frames with Ethernets 0x8890, 0x8891, and 0x8893 to pass.

Alternatively, you can use the following CLI options to change the Ethernets of the HA heartbeat packets:

```
config system ha
  set ha-eth-type <ha_ethertype_4-digit_hex>
  set hc-eth-type <hc_ethertype_4-digit_hex>
  set l2ep-eth-type <l2ep_ethertype_4-digit_hex>
end
```

For example, use the following command to change the Ethertype of the HA heartbeat packets from 0x8890 to 0x8895 and to change the Ethertype of HA Telnet session packets from 0x8891 to 0x889f:

```
config system ha
  set ha-eth-type 8895
  set l2ep-eth-type 889f
end
```

Enabling or disabling HA heartbeat encryption and authentication

You can enable HA heartbeat encryption and authentication to encrypt and authenticate HA heartbeat packets. HA heartbeat packets should be encrypted and authenticated if the cluster interfaces that send HA heartbeat packets are also connected to your networks.

If HA heartbeat packets are not encrypted the cluster password and changes to the cluster configuration could be exposed and an attacker may be able to sniff HA packets to get cluster information. Enabling HA heartbeat message authentication prevents an attacker from creating false HA heartbeat messages. False HA heartbeat messages could affect the stability of the cluster.

HA heartbeat encryption and authentication are disabled by default. Enabling HA encryption and authentication could reduce cluster performance. Use the following CLI command to enable HA heartbeat encryption and authentication.

```
config system ha
  set authentication enable
  set encryption enable
end
```

HA authentication and encryption uses AES-128 for encryption and SHA1 for authentication.

FGSP - FortiGate Session Life Support Protocol

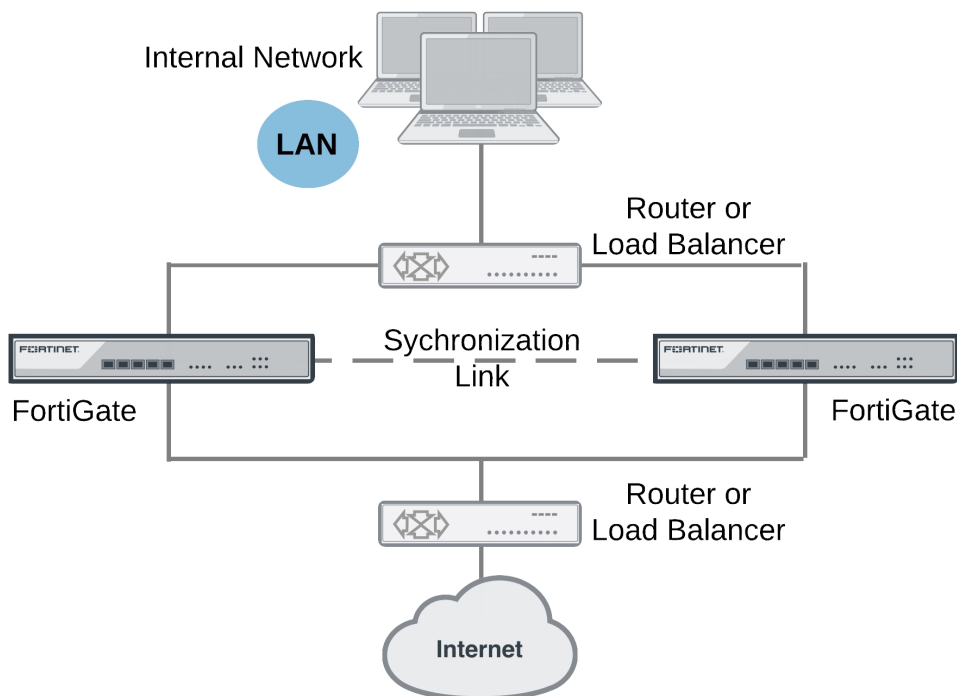
FortiGate Session Life Support Protocol (FGSP) distributes sessions between two FortiGate units and the FGSP performs session synchronization. If one of the peers fails, session failover occurs and active sessions fail over to the peer that is still operating. This failover occurs without any loss of data. Also, the external routers or load balancers will detect the failover and re-distribute all sessions to the peer that is still operating. The two FortiGate units must be the same model and must be running the same firmware.



Note that you cannot configure FGSP HA when FGCP HA is enabled.

You can also use the `config system cluster-sync` command to configure FGSP between two FortiGate units.

The FortiGate's HA Heartbeat listens on ports TCP/703, TCP/23, or ETH Layer 2/8890.



In previous versions of FortiOS, FGSP was called TCP session synchronization or standalone session synchronization. However, FGSP has been expanded to include both IPv4 and IPv6 TCP, UDP, ICMP, expectation, NAT sessions, and IPsec tunnels.

Configuration synchronization

Configuration synchronization can also be performed, allowing you to make configuration changes once for both FortiGate units instead of requiring multiple configuration changes on each FortiGate unit. However interface IP addresses, BGP neighbor settings, and other settings that identify the FortiGate unit on the network are not synchronized. You can enable configuration synchronization by entering the following command:

```
config system ha
    set standalone-config-sync enable
end
```

UDP and ICMP (connectionless) session synchronization

In many configurations, due to their non-stateful nature, UDP and ICMP sessions don't need to be synchronized to naturally failover. However, if it is required, you can configure the FGSP to also synchronize UDP and ICMP sessions by entering the following command:

```
config system ha
    set session-pickup enable
    set session-pickup-connectionless enable
end
```

Expectation (asymmetric) session synchronization

Synchronizing asymmetric traffic can be very useful in situations where multiple Internet connections from different ISPs are spread across two FortiGates.

The FGSP enforces firewall policies for asymmetric traffic, including cases where the TCP 3-way handshake is split between two FortiGates. For example, FGT-A receives the TCP-SYN, FGT-B receives the TCP-SYN-ACK, and FGT-A receives the TCP-ACK. Under normal conditions a firewall will drop this connection since the 3-way handshake was not seen by the same firewall. However two FortiGates with FGSP configured will be able to properly pass this traffic since the firewall sessions are synchronized.

If traffic will be highly asymmetric, as described above, the following command must be enabled on both FortiGates:

```
config system ha
    set session-pickup enable
    set session-pickup-expectation enable
end
```

Security profile inspection with asymmetric and symmetric traffic

Security profile inspection, flow or proxy based, is **not** expected to work properly if the traffic in the session is load balanced across more than one FortiGate in either direction. However, flow-based inspection should be used in FGSP deployments.

For symmetric traffic, security profile inspection can be used but with the following limitations:

- No session synchronization for the sessions inspected using proxy-based inspection. Sessions will drop and need to be reestablished after data path failover.
- Sessions with flow-based inspection will failover, and inspection of sessions after a failover may not work.

Improving session synchronization performance

Two HA configuration options are available to reduce the performance impact of enabling session failover (also known as session pickup): reducing the number of sessions that are synchronized by adding a session pickup delay, and using more FortiGate interfaces for session synchronization.

Reducing the number of sessions that are synchronized

If session pickup is enabled, as soon as new sessions are added to the primary unit session table they are synchronized to the other cluster units. Enable the session-pickup-delay CLI option to reduce the number of sessions that are synchronized by synchronizing sessions only if they remain active for more than 30 seconds. Enabling this option could greatly reduce the number of sessions that are synchronized if a cluster typically processes very many short duration sessions, which is typical of most HTTP traffic for example.

Use the following command to enable a 30 second session pickup delay:

```
config system ha
    set session-pickup-delay enable
end
```

Enabling session pickup delay means that if a failover occurs more sessions may not be resumed after a failover. In most cases short duration sessions can be restarted with only a minor traffic interruption. However, if you notice too many sessions not resuming after a failover you might want to disable this setting.

Using multiple FortiGate interfaces for session synchronization

Using the `session-sync-dev` option, you can select one or more FortiGate interfaces to use for synchronizing sessions as required for session pickup. Normally session synchronization occurs over the HA heartbeat link. Using this HA option means only the selected interfaces are used for session synchronization and not the HA heartbeat link. If you select more than one interface, session synchronization traffic is load balanced among the selected interfaces.

Moving session synchronization from the HA heartbeat interface reduces the bandwidth required for HA heartbeat traffic and may improve the efficiency and performance of the cluster, especially if the cluster is synchronizing a large number of sessions. Load balancing session synchronization among multiple interfaces can further improve performance and efficiency if the cluster is synchronizing a large number of sessions.

Use the following command to perform cluster session synchronization using the port10 and port12 interfaces.

```
config system ha
    set session-sync-dev port10 port12
end
```

Session synchronization packets use Ethertype 0x8892. The interfaces to use for session synchronization must be connected together either directly using the appropriate cable (possible if there are only two units in the cluster) or using switches. If one of the interfaces becomes disconnected the cluster uses the remaining interfaces for session synchronization. If all of the session synchronization interfaces become disconnected, session synchronization reverts back to using the HA heartbeat link. All session synchronization traffic is between the primary unit and each subordinate unit.

Since large amounts of session synchronization traffic can increase network congestion, it is recommended that you keep this traffic off of your network by using dedicated connections for it.



Note that "unsetting" `session-sync-dev` (i.e. by entering `unset session-sync-dev`) has the following two effects:

1. Session synchronization will use the ports defined as HA heartbeat interfaces (`set hbdev`).
2. Session synchronization packets will be sent over UDP/708 instead of Ethertype 0x8892.

NAT session synchronization

NAT sessions are not synchronized by default. You can enable NAT session synchronization by entering the following command:

```
config system ha
    set session-pickup enable
    set session-pickup-nat enable
end
```

Note that, after a failover with this configuration, all sessions that include the IP addresses of interfaces on the failed FortiGate unit will have nowhere to go since the IP addresses of the failed FortiGate unit will no longer be on the network. If you want NAT sessions to resume after a failover you should not configure NAT to use the destination interface IP address, since the FGSP FortiGate units have different IP addresses. To avoid this issue, you should use IP pools with the type set to `overload` (which is the default IP pool type), as shown in the example below:

```
config firewall ippool
    edit FGSP-pool
        set type overload
        set startip 172.20.120.10
        set endip 172.20.120.20
    end
```

In NAT/Route mode, only sessions for route mode security policies are synchronized. FGSP HA is also available for FortiGate units or virtual domains operating in Transparent mode. Only sessions for normal Transparent mode policies are synchronized.

IPsec tunnel synchronization

When you use the `config system cluster-sync` command to enable FGSP, IPsec keys and other runtime data are synchronized between cluster units. This means that if one of the cluster units goes down the cluster unit that is still operating can quickly get IPsec tunnels re-established without re-negotiating them. However, after a failover, all existing tunnel sessions on the failed FortiGate have to be restarted on the still operating FortiGate.

IPsec tunnel sync only supports dialup IPsec. The interfaces on both FortiGates that are tunnel endpoints must have the same IP addresses and external routers must be configured to load balance IPsec tunnel sessions to the FortiGates in the cluster.

Standalone configuration synchronization uses a very similar process as FGCP. There is a similar relationship between the two FortiGates but only in regards to configuration synchronization, not session information. The primary unit is selected by using priority/override. The heartbeat is used to check the primary unit's health. Once heartbeat loss is detected, a new primary unit is selected.

Automatic session synchronization after peer reboot

The following command allows you to configure an automatic session synchronization after a peer FGSP unit has rebooted. FGSP will send out heartbeat signals (every 1 - 10 seconds, as shown below) if one FortiGate is rebooting and the other FortiGate fails.

To configure automatic session synchronization:

```
config system session-sync
  edit 1
    set down-intfs-before-sess-sync <interfaces> - List of interfaces to be turned down before session
      synchronization is complete.
    set-hb-interval <integer> - (1 - 10 seconds)
    set hb-lost-threshold <integer> - (1 - 10)
  next
end
```

FGFM - FortiGate to FortiManager Protocol

The FortiGate to FortiManager (FGFM) protocol is designed for FortiGate and FortiManager deployment scenarios, especially where NAT is used. These scenarios include the FortiManager on public internet while the FortiGate unit is behind NAT, FortiGate unit is on public internet while FortiManager is behind NAT, or both FortiManager and FortiGate unit have routable IP addresses.

The FortiManager unit's Device Manager uses FGFM to create new device groups, provision and add devices, and install policy packages and device settings.

Port 541 is the default port used for FortiManager traffic on the internal management network.

Adding a FortiGate to the FortiManager

Adding a FortiGate unit to a FortiManager requires configuration on both devices. This section describes the basics to configure management using a FortiManager device.

FortiGate configuration

Adding a FortiGate unit to FortiManager will ensure that the unit will be able to receive antivirus and IPS updates and allow remote management through the FortiManager system, or FortiCloud service. The FortiGate unit can be in either NAT or transparent mode. The FortiManager unit provides remote management of a FortiGate unit over TCP port 541.

You must first enable **Central Management** on the FortiGate so management updates to firmware and FortiGuard services are available:

1. Go to **Security Fabric > Settings**.
2. Enable **Central Management** and set **Type** to **FortiManager**.
3. Enter the FortiManager's **IP/Domain Name** in the field provided.

To configure the previous steps in the CLI, enter the following - note that `fmg` can be set to either an IP address or FQDN:

```
config system central-management
    set fmg <string>
end
```

To use the registration password, enter the following:

```
execute central-mgmt register-device <fmg-serial-no> <fmg-register-password> <fgtusrname>
    <fgt-password>
```

FGFM is also used in ADOMs (Administrative Domains) set to Normal Mode. Normal Mode has Read/Write privileges, where the administrator is able to make changes to the ADOM and manage devices from the FortiManager. FortiGate units in the ADOM will query their own configuration every five seconds. If there has been a configuration change, the FortiGate unit will send a revision on the change to the FortiManager using the FGFM protocol.

To configure central management on the FortiGate unit, enter the following on the FortiGate:

```
config system central-management
    set mode normal
    set fortimanager-fds-override enable
    set fmg <string>
```


end

Configuring an SSL connection

The default encryption automatically sets high and medium encryption algorithms. Algorithms used for **High**, **Medium**, and **Low** follow the openssl definitions below:

Encryption level	Key strength	Algorithms used
High	Key lengths larger than 128 bits, and some cipher suites with 128-bit keys.	DHE-RSA-AES256-SHA:AES256-SHA: EDH-RSA-DES-CBC3-SHA: DES-CBC3-SHA:DES-CBC3-MD5:DHE-RSA-AES128-SHA:AES128-SHA
Medium	Key strengths of 128 bit encryption.	RC4-SHA:RC4-MD5:RC4-MD
Low	Key strengths of 64 or 56 bit encryption algorithms but excluding export cipher suites.	EDH-RSA-DES-CBC-SHA; DES-CBC-SHA; DES-CBC-MD5

An SSL connection can be configured between the two devices and an encryption level selected. To configure the connection in the CLI, enter the following:

```
config system central-management
  set status enable
  set enc-algorithm (default | high | low)
end
```

Note that `default` automatically sets high and medium encryption algorithms.

FortiManager configuration

Use the **Device Manager** pane to add, configure, and manage devices.

You can add existing operational devices, unregistered devices, provision new devices, and add multiple devices at a time.

Adding an operating FortiGate HA cluster to the **Device Manager** pane is similar to adding a standalone device. Type the IP address of the master device. The FortiManager will handle the cluster as a single managed device.



To confirm that a device model or firmware version is supported by current firmware version running on FortiManager, enter the following CLI command:

```
diagnose dvm supported-platforms list
```

See the [FortiManager Administration Guide](#) for full details on adding devices, under **Device Manager**.

Replacing a FortiGate in a FortiManager configuration

FGFM can be used in order to re-establish a connection between a FortiGate unit and a FortiManager configuration. This is useful for if you need a FortiGate unit replaced following an RMA hardware replacement. This applies to a FortiGate running in HA as the primary units; it does not apply to subordinate units.

When the FortiGate unit is replaced, perform a Device Manager Connectivity check or Refresh on the FortiManager to establish the FGFM management tunnel to the FortiGate. If it fails to establish, you can force the tunnel by executing the following command on the FortiManager:

```
execute fgfm reclaim-dev-tunnel <device_name>
```

Debugging FGFM on FortiManager

- To display diagnostic information for troubleshooting, set the debug level of the FGFM daemon (enter a device name to only show messages related to that device):

```
diagnose debug application fgfmsd <integer> <device_name>
```

- To view installation session, object, and session lists:

```
diagnose fgfm install-session
diagnose fgfm object-list
diagnose fgfm session-list <device_ID>
```

- To reclaim a management tunnel (device name is optional):

```
execute fgfm reclaim-dev-tunnel <device_name>
```

- To view the link-local address assigned to the FortiManager:

```
diagnose fmnetwork interface list
```

Debugging FGFM on FortiGate

- To view information about the Central Management System configuration:

```
get system central-management
```

- To produce realtime debugging information:

```
diagnose debug application fgfmd -1
```

- To view the link-local address assigned to the FortiManager:

```
diagnose fmnetwork interface list
```

FortiOS DHCP options and auto DNS hostname for FortiManager details

A diagnose command can be used to show the FortiManager autodiscovery status for the secure sending of FortiManager details to FortiGate.

FortiGate is occasionally required in large deployments where a Zero Touch Provisioning (ZTP) of the unit is required.

Rather than using the CLI Console to configure system settings one at a time, ZTP can help to reduce errors, save time in automated device configuration, and enhance scalability.

This functionality is designed to work even in a closed network with no Internet access.

To verify the FortiManager autodiscovery status, use the following command:

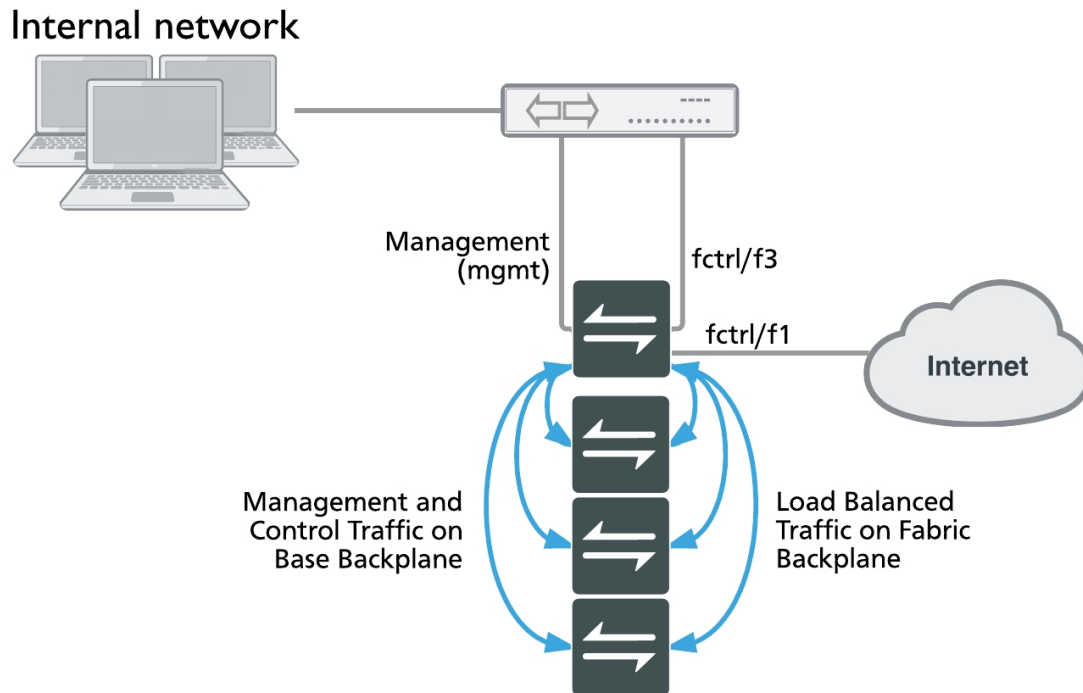
```
diagnose fdsm fmg-auto-discovery-status
```

SLBC - Session-aware Load Balancing Cluster

The Session-aware Load Balancing Cluster (SLBC) protocol is used for clusters consisting of FortiControllers that perform load balancing of both TCP and UDP sessions. As session-aware load balancers, FortiControllers, with FortiASIC DP processors, are capable of directing any TCP or UDP session to any worker installed in the same chassis. It also means that more complex networking features such as NAT, fragmented packets, complex UDP protocols and others such as Session Initiation Protocol (SIP), a communications protocol for signaling and controlling multimedia communication sessions, can be load balanced by the cluster.

Currently, only three FortiController models are available for SLBC: FortiController-5103B, FortiController-5903C, and FortiController-5913C. Supported workers include the FortiGate-5001B, 5001C, 5101C, and 5001D.

FortiGate-7000 series products also support SLBC.



You cannot mix FGCP and SLBC clusters in the same chassis.

An SLBC with two FortiControllers can operate in active-passive mode or dual mode. In active-passive mode, if the active FortiController fails, traffic is transferred to the backup FortiController. In dual mode both FortiControllers load balance traffic and twice as many network interfaces are available.

SLBC clusters consisting of more than one FortiController use the following types of communication between FortiControllers to operate normally:

- **Heartbeat:** Allows the FortiControllers in the cluster to find each other and share status information. If a FortiController stops sending heartbeat packets it is considered down by other cluster members. By default heartbeat traffic uses VLAN 999.
- **Base control:** Communication between FortiControllers on subnet 10.101.11.0/255.255.255.0 using VLAN 301.
- **Base management:** Communication between FortiControllers on subnet 10.101.10.0/255.255.255.0 using VLAN 101.
- **Session synchronization:** If one FortiController fails, session synchronization allows another to take its place and maintain active communication sessions. FortiController-5103B session sync traffic uses VLAN 2000. FortiController-5903C and FortiController-5913C session sync traffic between the FortiControllers in slot 1 uses VLAN 1900 and between the FortiControllers in slot 2 uses VLAN 1901. You cannot change these VLANs.

Note that SLBC does not support session synchronization between workers in the same chassis. The FortiControllers in a cluster keep track of the status of the workers in their chassis and load balance sessions to the workers. If a worker fails the FortiController detects the failure and stops load balancing sessions to that worker. The sessions that the worker is processing when it fails are lost.

Changing the heartbeat VLAN

To change the VLAN from the FortiController GUI, from the **System Information** dashboard widget, beside **HA Status**, select **Configure**. Change the **VLAN to use for HA heartbeat traffic(1-4094)** setting.

You can also change the heartbeat VLAN ID from the FortiController CLI. For example, to change the heartbeat VLAN ID to **333**, enter the following:

```
config system ha
    set hbdev-vlan-id 333
end
```

Setting the mgmt interface as a heartbeat interface

To add the mgmt interface to the list of heartbeat interfaces used, on the FortiController-5103B, enter the following:

```
config system ha
    set hbdev b1 b2 mgmt
end
```

This example adds the mgmt interface for heartbeats to the B1 and B2 interfaces. The B1 and B2 ports are recommended because they are 10G ports and the mgmt interface is a 100Mbit interface.



FortiController-5103B is currently the only model that allows its mgmt interface to be added to the heartbeat interfaces list.

Changing the heartbeat interface mode

By default, only the first heartbeat interface (usually B1) is used for heartbeat traffic. If this interface fails on any of the FortiControllers in a cluster, then the second heartbeat interface is used (B2).

To simultaneously use all heartbeat interfaces for heartbeat traffic, enter the following command:

```
config load-balance-setting
    set base-mgmt-interface-mode active-active
end
```

Changing the base control subnet and VLAN

You can change the base control subnet and VLAN from the FortiController CLI. For example, to change the base control subnet to **10.122.11.0/255.255.255.0** and the VLAN ID to **320**, enter the following:

```
config load-balance setting
  set base-ctrl-network 10.122.11.0 255.255.255.0
  config base-ctrl-interfaces
    edit b1
      set vlan-id 320
    next
    edit b2
      set vlan-id 320
  end
```

Changing the base management subnet and VLAN

You can change the base management subnet from the FortiController GUI under **Load Balance > Config** and changing the **Internal Management Network**.

You can also change the base management subnet and VLAN ID from the FortiController CLI. For example, to change the base management subnet to **10.121.10.0/255.255.255.0** and the VLAN to **131**, enter the following:

```
config load-balance setting
  set base-mgmt-internal-network 10.121.10.0 255.255.255.0
  config base-mgt-interfaces
    edit b1
      set vlan-id 131
    next
    edit b2
      set vlan-id 131
  end
```

If required, you can use different VLAN IDs for the B1 and B2 interface.

Changing this VLAN only changes the VLAN used for base management traffic between chassis. Within a chassis the default VLAN is used.

Enabling and configuring the session sync interface

To enable session synchronization in a two chassis configuration, enter the following command:

```
config load-balance setting
  set session-sync enable
end
```

You will then need to select the interface to use for session sync traffic. The following example sets the FortiController-5103B session sync interface to **F4**:

```
config system ha
  set session-sync-port f4
end
```

The FortiController-5903C and FortiController-5913C use B1 and B2 as the session sync interfaces so no configuration changes are required.

FGCP to SLBC migration

You can convert an FGCP virtual cluster (with VDOMs) to an SLBC cluster. The conversion involves replicating the VDOM, interface, and VLAN configuration of the FGCP cluster on the SLBC cluster primary worker, then backing up the configuration of each FGCP cluster VDOM. Each of the VDOM configuration files is manually edited to adjust interface names. These modified VDOM configuration files are then restored to the corresponding SLBC cluster primary worker VDOMs.

For this migration to work, the FGCP cluster and the SLBC workers must be running the same firmware version, the VDOMs are enabled on the FGCP cluster, and the SLBC workers have been registered and licensed. However, the FGCP cluster units do not have to be the same model as the SLBC cluster workers.

Only VDOM configurations are migrated. You have to manually configure primary worker management and global settings.

Conversion steps

1. Add VDOM(s) to the SLBC primary worker with names that match those of the FGCP cluster.
2. Map FGCP cluster interface names to SLBC primary worker interface names. For example, you can map the FGCP cluster port1 and port2 interfaces to the SLBC primary worker fctl/f1 and fctl/f2 interfaces. You can also map FGCP cluster interfaces to SLBC trunks, and include aggregate interfaces.
3. Add interfaces to the SLBC primary worker VDOMs according to your mapping. This includes moving SLBC physical interfaces into the appropriate VDOMs, creating aggregate interfaces, and creating SLBC trunks if required.
4. Add VLANs to the SLBC primary worker that match VLANs in the FGCP cluster. They should have the same names as the FGCP VLANs, be added to the corresponding SLBC VDOMs and interfaces, and have the same VLAN IDs.
5. Add inter-VDOM links to the SLBC primary worker that match the FGCP cluster.
6. Backup the configurations of each FGCP cluster VDOM, and SLBC primary worker VDOM.
7. Use a text editor to replace the first four lines of each FGCP cluster VDOM configuration file with the first four lines of the corresponding SLBC primary worker VDOM configuration file. Here are example lines from an SLBC primary worker VDOM configuration file:


```
#config-version=FG-5KB-5.02-FW-build670-150318:opmode=0:vdom=1:user=admin
#conf_file_ver=2306222306838080295
#buildno=0670
#global_vdom=0:vd_name=VDOM1
```
8. With the text editor, edit each FGCP cluster VDOM configuration file and replace all FGCP cluster interface names with the corresponding SLBC worker interface names, according to the mapping you created in step 2.
9. Set up a console connection to the SLBC primary worker to check for errors during the following steps.
10. From the SLBC primary worker, restore each FGCP cluster VDOM configuration file to each corresponding SLBC primary worker VDOM.
11. Check the following on the SLBC primary worker:
 - Make sure `set type fctrl-trunk` is enabled for SLBC trunk interfaces.
 - Enable the global and management VDOM features that you need, including SNMP, logging, connections to FortiManager, FortiAnalyzer, and so on.
 - If there is a FortiController in chassis slot 2, make sure the worker base2 interface status is up.
 - Remove `snmp-index` entries for each interface.

- Since you can manage the workers from the FortiController you can remove management-related configurations using the worker mgmt1 and mgmt2 interfaces (Logging, SNMP, admin access, etc.) if you are not going to use these interfaces for management.

How to set up SLBC with one FortiController-5103B

This example describes the basics of setting up a Session-aware Load Balancing Cluster (SLBC) that consists of one FortiController-5103B, installed in chassis slot 1, and three FortiGate-5001C workers, installed in chassis slots 3, 4, and 5.

This SLBC configuration can have up to eight 10Gbit network connections.

Configuring the hardware

1. Install a FortiGate-5000 series chassis and connect it to power. Install the FortiController in slot 1. Install the workers in slots 3, 4, and 5. Power on the chassis.
2. Check the chassis, FortiController, and FortiGate LEDs to verify that all components are operating normally. (To check normal operation LED status see the FortiGate-5000 series documents available [here](#).)
3. Check the FortiSwitch-ATCA release notes and install the latest supported firmware on the FortiController and on the workers. Get FortiController firmware from the [Fortinet Support site](#). Select the FortiSwitch-ATCA product.

Configuring the FortiController

To configure the FortiController, you will need to either connect to the FortiController GUI or CLI with the default IP address of `http://192.168.1.99`. Log in using the admin account (no password).

1. Add a password for the admin account. Use the **Administrators** widget in the GUI, or enter the following CLI command:

```
config admin user
  edit admin
    set password <password>
  end
```
2. Change the FortiController mgmt interface IP address. Use the **Management Port** widget in the GUI, or enter the following CLI command:

```
config system interface
  edit mgmt
    set ip 172.20.120.151/24
  end
```
3. If you need to add a default route for the management IP address, enter the following command:

```
config route static
  edit route 1
    set gateway 172.20.121.2
  end
```
4. To set the chassis type that you are using, enter the following CLI command:

```
config system global
  set chassis-type fortigate-5140
end
```
5. Go to **Load Balance > Config** and add workers to the cluster by selecting **Edit** and moving the slots that contain workers to the **Member** list. The Config page shows the slots in which the cluster expects to find workers. Since the workers have not been configured yet, their status is **Down**.
Configure the **External Management IP/Netmask**. Once the workers are connected to the cluster, you can use

this IP address to manage and configure them.

6. You can also enter the following CLI command to add slots 3, 4, and 5 to the cluster:

```
config load-balance setting
  config slots
    edit 3
  next
  edit 4
  next
  edit 5
  end
end
```

7. You can also enter the following command to configure the external management IP/Netmask and management access to the following address:

```
config load-balance setting
  set base-mgmt-external-ip 172.20.120.100 255.255.255.0
  set base-mgmt-allowaccess https ssh ping
end
```

Adding the workers

Before you begin adding workers to the cluster, make sure you enter the `execute factoryreset` command in the CLI so the workers are set to factory default settings. If the workers are going to run FortiOS Carrier, add the FortiOS Carrier licence instead - this will reset the worker to factory default settings.

Also make sure to register and apply licenses to each worker, including FortiClient licensing, FortiCloud activation, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMs). You can also install any third-party certificates on the primary worker before forming the cluster. Once the cluster is formed, third-party certificates are synchronized to all of the workers. FortiToken licenses can be added at any time, which will also synchronize across all of the workers.

1. Log in to each of the worker's CLI and enter the following CLI command to set the worker to operate in FortiController mode:

```
config system elbc
  set mode fortincontroller
end
```

Once the command is entered, the worker restarts and joins the cluster.

2. On the FortiController, go to **Load Balance > Status**. You will see the workers appear in their appropriate slots. The worker in the lowest slot number usually becomes the primary unit.

You can now manage the workers in the same way as you would manage a standalone FortiGate. You can connect to the worker GUI or CLI using the **External Management IP**. If you had configured the worker mgmt1 or mgmt2 interfaces you can also connect to one of these addresses to manage the cluster.

To operate the cluster, connect networks to the FortiController front panel interfaces and connect to a worker GUI or CLI to configure the workers to process the traffic they receive. When you connect to the External Management IP you connect to the primary worker. When you make configuration changes they are synchronized to all workers in the cluster.

Managing the devices in an SLBC with the external management IP

The External Management IP address is used to manage all of the individual devices in a SLBC by adding a special port number. This special port number begins with the standard port number for the protocol you are using

and is followed by two digits that identify the chassis number and slot number. The port number can be calculated using the following formula:

$$\text{service_port} \times 100 + (\text{chassis_id} - 1) \times 20 + \text{slot_id}$$

Where:

- **service_port** is the normal port number for the management service (80 for HTTP, 443 for HTTPS and so on).
- **chassis_id** is the chassis ID specified as part of the FortiController HA configuration and can be 1 or 2.
- **slot_id** is the number of the chassis slot.



By default, chassis 1 is the primary chassis and chassis 2 is the backup chassis. However, the actual primary chassis is the one with the primary FortiController, which can be changed independently of the chassis number. Additionally, the **chassis_id** is defined by the chassis number, *not* whether the chassis contains the primary FortiController.

Some examples:

- HTTPS, chassis 1, slot 2: $443 \times 100 + (1 - 1) \times 20 + 2 = 44300 + 0 + 2 = 44302$:
browse to: <https://172.20.120.100:44302>
- HTTP, chassis 2, slot 4: $80 \times 100 + (2 - 1) \times 20 + 4 = 8000 + 20 + 4 = 8024$:
browse to <http://172.20.120.100/8024>
- HTTPS, chassis 1, slot 10: $443 \times 100 + (1 - 1) \times 20 + 10 = 44300 + 0 + 10 = 44310$:
browse to <https://172.20.120.100/44310>

Single chassis or chassis 1 special management port numbers

Slot number	HTTP (80)	HTTPS (443)	Telnet (23)	SSH (22)	SNMP (161)
Slot 1	8001	44301	2301	2201	16101
Slot 2	8002	44302	2302	2202	16102
Slot 3	8003	44303	2303	2203	16103
Slot 4	8004	44304	2304	2204	16104
Slot 5	8005	44305	2305	2205	16105
Slot 6	8006	44306	2306	2206	16106
Slot 7	8007	44307	2307	2207	16107
Slot 8	8008	44308	2308	2208	16108
Slot 9	8009	44309	2309	2209	16109

Slot number	HTTP (80)	HTTPS (443)	Telnet (23)	SSH (22)	SNMP (161)
Slot 10	8010	44310	2310	2210	16110
Slot 11	8011	44311	2311	2211	16111
Slot 12	8012	44312	2312	2212	16112
Slot 13	8013	44313	2313	2213	16113
Slot 14	8014	44314	2314	2214	16114

Chassis 2 special management port numbers


Slot number	HTTP (80)	HTTPS (443)	Telnet (23)	SSH (22)	SNMP (161)
Slot 1	8021	44321	2321	2221	16121
Slot 2	8022	44322	2322	2222	16122
Slot 3	8023	44323	2323	2223	16123
Slot 4	8024	44324	2324	2224	16124
Slot 5	8025	44325	2325	2225	16125
Slot 6	8026	44326	2326	2226	16126
Slot 7	8027	44327	2327	2227	16127
Slot 8	8028	44328	2328	2228	16128
Slot 9	8029	44329	2329	2229	16129
Slot 10	8030	44330	2330	2230	16130
Slot 11	8031	44331	2331	2231	16131
Slot 12	8032	44332	2332	2232	16132
Slot 13	8033	44333	2333	2233	16133
Slot 14	8034	44334	2334	2234	16134

For more detailed information regarding FortiController SLBC configurations, see the [FortiController Session-Aware Load Balancing \(SLBC\) Guide](#).

Fortinet Security Fabric

The Fortinet Security Fabric spans across an entire network linking different security sensors and tools together to collect, coordinate, and respond to malicious behavior in real time. Security Fabric can be used to coordinate the behavior of different Fortinet products in your network, including FortiGate, FortiAnalyzer, FortiClient, FortiSandbox, FortiAP, FortiSwitch, and FortiClient Enterprise Management Server (EMS). Security Fabric supports FortiOS 5.4.1+, FortiSwitchOS 3.3+, and FortiClient 5.4.1+.

Port TCP/8009 is the port FortiGate uses for incoming traffic from the FortiClient Portal, as user information (such as IP address, MAC address, avatar, and other profile information) is automatically synchronized to the FortiGate and EMS.



The brief example below assumes that FortiTelemetry has been enabled on the top-level FortiGate (**FGT1**), OSPF routing has been configured, and that policies have been created for all FortiGate units to access the Internet.

For more details on how to configure a security fabric between FortiGate units, see [Fortinet Security Fabric installation](#) on the Fortinet Cookbook website.


Enabling Security Fabric on the FortiGate:

1. On the upstream FortiGate (FGT1), go to **Security Fabric > Settings** and enable **FortiGate Telemetry**.
2. Enter a **Group name** and **Group password** for the fabric.

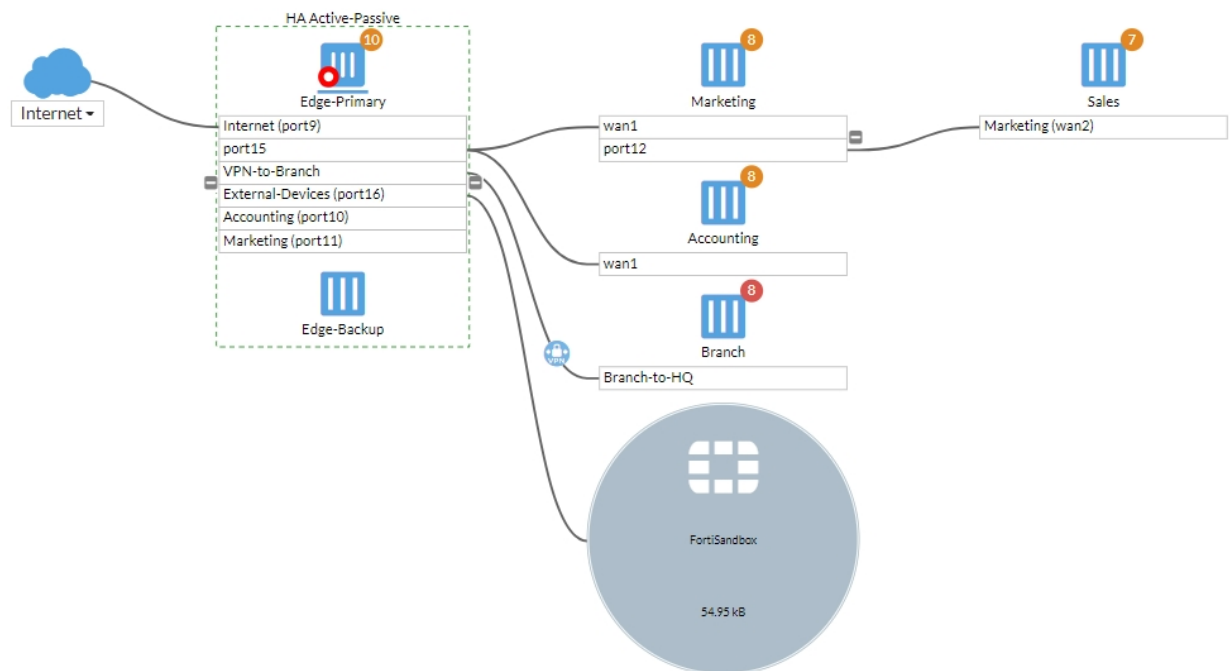
3. On a downstream FortiGate (such as FGT2 or FGT3), configure the same fabric settings as were set on FGT1.
4. Enable **Connect to upstream FortiGate**.
Be sure you do *not* enable this on the topmost-level FortiGate (in this example, FGT1).
5. In **FortiGate IP**, enter the FGT1 interface that has **FortiTelemetry** enabled. The **Management IP** can be left to use the WAN IP or optionally specified.

Once set up, you can view your network's Security Fabric configuration under **Security Fabric** through two topology dashboards.

6. On the top-level FortiGate, go to **Security Fabric > Physical Topology**. This dashboard shows a visualization of all access layer devices in the fabric.



7. Go to **Security Fabric > Logical Topology** to view information about the interfaces (logical or physical) that each device in the fabric is connected to.



Other Security Fabric configurations for your network are available through the Fortinet Cookbook [Security Fabric Collection](#) page.

FortiTelemetry/On-Net/FortiClient Endpoint Compliance

FortiTelemetry (called FortiHeartBeat in FortiOS 5.4.0 and FortiClient Access in FortiOS 5.2) is an interface option that listens for connections from devices with FortiClient installed.

FortiTelemetry is the TCP/8013 protocol used between FortiClient and FortiGate, FortiClient and FortiClient EMS, and between FortiGate and other FortiGates in CSF configurations.



While all GUI references of FortiHeartBeat have been changed to FortiTelemetry in FortiOS 5.4.1, the CLI options have *not* been renamed and will remain as `fortiheartbeat`.

With FortiTelemetry enabled on the FortiGate, you can enforce FortiTelemetry for all FortiClients. This FortiClient endpoint compliance will require all clients to have FortiClient installed in order to get access through the FortiGate. Configure these settings in the internal interface under **Network > Interfaces**. Edit the interface of your choice. Under **Administrative Access**, enable **FortiTelemetry**, then enable **FortiClient On-Net Status**.

To enable FortiTelemetry on an interface - CLI:

```
config system interface edit <port_number>
    set listen-forticlient-connection enable
    set endpoint-compliance enable
end
```

You can also enable **DHCP server** and **FortiClient On-Net Status** to display the on-net status of FortiClient devices on the FortiClient Monitor (under **Monitor > FortiClient Monitor**).

To enable FortiClient On-Net status for a DHCP server added to the port1 interface - CLI:

```
config system dhcp server edit 1
    set interface port1
    set forticlient-on-net-status enable
end
```

FortiClient endpoint licence updates

FortiClient endpoint licenses for FortiOS 5.6.0 can be purchased in multiples of 100. There is a maximum client limit based on the FortiGate's model. FortiCare enforces the maximum limits when the customer is applying the license to a model.

If you are using the ten free licenses for FortiClient, support is provided on the Fortinet Forum (forum.fortinet.com). Phone support is only available for paid licenses.

Model(s)	Maximum Client Limit
VM00	200
FGT/FWF 30 to 90 series	200
FGT 100 to 400 series	600
FGT 500 to 900 series, VM01, VM02	2,000
FGT 1000 to 2900 series	20,000
FGT 3000 to 3600 series, VM04	50,000
FGT 3700D and above, VM08 and above	100,000

Older FortiClient SKUs will still be valid and can be applied to FortiOS 5.4 and 5.6.

Connecting FortiClient Telemetry after installation

After FortiClient is installed on an endpoint, FortiClient automatically launches and searches for a FortiGate or FortiClient EMS for FortiClient Telemetry connection. When FortiClient locates a FortiGate or EMS, the **FortiGate Detected** or **Enterprise Management Server (EMS) Detected** dialog box will appear:



If all the information displayed is correct, select **Accept**. FortiClient Telemetry will connect to the identified FortiGate/EMS.

Alternately, you can select **Cancel** and launch FortiClient without connecting to FortiClient Telemetry. This will launch FortiClient in standalone mode, where you can manually connect FortiClient Telemetry.

After FortiClient Telemetry is connected to FortiGate or EMS, FortiClient downloads a profile from FortiGate/EMS.

How FortiClient locates FortiGate/EMS

FortiClient uses the following methods in the following order to automatically locate FortiGate/EMS for Telemetry connection:

- 1. Telemetry gateway IP list:** FortiClient Telemetry searches for IP addresses in its subnet in the Gateway IP list. It connects to the FortiGate in the list that is also in the same subnet as the host system.
If FortiClient cannot find any FortiGates in its subnet, it will attempt to connect to the first reachable FortiGate in the list, starting from the top. The order of the list is maintained as it was configured in the Gateway IP list.
- 2. Remembered gateway IP list:** You can configure FortiClient to remember gateway IP addresses when you connect Telemetry to FortiGate/EMS. Later FortiClient can use the remembered IP addresses to automatically connect Telemetry to FortiGate/EMS.
- 3. Default gateway IP address:** The default gateway IP address is specified on the FortiClient endpoint and is used to automatically connect to FortiGate. This method does not support connection to EMS.



FortiClient obtains the default gateway IP address from the operating system on the endpoint device. The default gateway IP address of the endpoint device should be the IP address for the FortiGate interface with Telemetry enabled.

If FortiClient is unable to automatically locate a FortiGate/EMS on the network for Telemetry connection, you can type the gateway IP address of the FortiGate/EMS.



FortiClient uses the same process to connect Telemetry to FortiGate/EMS after the FortiClient endpoint reboots, rejoins the network, or encounters a network change.

FortiGuard

FortiGuard services can be purchased and registered to your FortiGate unit. The FortiGate must be connected to the Internet in order to automatically connect to the FortiGuard Distribution Network (FDN) to validate the license and download FDN updates.

The FortiGuard subscription update services include:

- AntiVirus (AV)
- Intrusion Protection Service (IPS)
- Application Control
- Anti-Spam
- Web Filtering
- Web Application Firewall (WAF)

The FDN sends notice that a FortiGuard AntiVirus and IPS update is available on UDP/9443.

The following information concerns certain considerations in regards to FortiGate receiving FortiGuard updates through FDN, how the submission of malware statistics to FortiGuard is handled, an automatic update behaviour when FortiGate has expired licenses, and related CLI commands.

Enabling FDN updates and FortiGuard services

In order to receive FortiGuard subscription updates, the unit needs to have access to the Internet and be able to connect to a DNS server in order to resolve the following URLs:

- **update.fortiguard.net:** For AV and IPS updates.
 - **service.fortiguard.net:** For web filtering and anti-spam updates.
1. Go to **System > FortiGuard**. Under **AntiVirus & IPS Updates**, enable **Scheduled Updates**, and configure an update schedule.
 2. You can force the unit to connect to the AV/IPS server by selecting **Update AV & IPS Definitions**.
 3. You can view your subscription details above in the **License Information** table.
 4. Once the schedule has been enabled, select **Apply**.

To see if the service is viable, open the CLI console and enter the following commands below.

For Web Filtering:

```
diagnose debug rating
```

For Anti-Spam:

```
diagnose spamfilter fortishield servers
```

If only one or two IPs are displayed in the command outputs, it could be one of the following issues:

- **No response from the DNS server:** Either the DNS server is unreachable or there is a problem with the routing. Make sure that contact to the DNS server is available by resolving some URLs from the CLI, for example:

```
execute ping www.google.com
execute ping service.fortiguard.net
```

You can also

- **Review update errors:** Review update information from the last update, enable debug outputs and force the update:

```
diagnose test update info
diagnose debug enable
diagnose debug application update 255
execute update-ase
execute update-av
execute update-ips
```

After troubleshooting, it is highly recommended to turn off debug mode:

```
diagnose debug disable
diagnose debug application update 0
```

- **FortiGuard Web filtering:** Port blocking or packet inspection is occurring downstream. The default port used by the FortiGuard for the FortiGuard services is 8888.

You can change this port using the following command:

```
config system fortiguard
    set port <port_number>
end
```

You can also change the source port for management traffic with the following CLI command:

```
config system global
    set ip-src-port-range 1035-25000
end
diagnose test application urlfilter 99
diagnose test application smtp 99
```

Submission of malware statistics to FortiGuard

FortiGates periodically send encrypted AntiVirus, IPS, botnet IP list, and Application Control event statistics to FortiGuard. Included with these malware statistics is the IP address and serial number of the FortiGate and the country in which the FortiGate is located.

The statistics are used to improve various aspects of FortiGate malware protection. For example, AntiVirus statistics allow FortiGuard to determine the viruses that are active in the wild. Signatures for such viruses are kept in the Active AV Signature Database that is used by many Fortinet products. Signatures for inactive viruses are moved to the Extended/Extreme AV Signature Database used by some customers. If the events for inactive viruses start appearing in malware statistics, these signatures can be moved back to the Active AV Signature Database.

The FortiGate and FortiGuard servers go through a 2-way SSL/TLS 1.2 authentication before any data is transmitted. The certificates used in this process must be trusted by each other and signed by Fortinet CA server.

Malware statistics are accumulated and sent periodically (by default every 60 minutes).

Fortinet products can only accept data from authorized FortiGuard servers. Fortinet products use DNS to find FortiGuard servers and periodically update their FortiGate server list. All other servers are provided by a list that is updated through the encrypted channel.



The submission of malware data is in accordance with Fortinet's "Automatically-Collected Information" detailed in the [Fortinet Privacy Policy](#), and the purpose of this collection is outlined in the "Use of your Information" section of the privacy policy.

There is no sensitive or personal information included in these submissions. Only malware statistics are sent.

Fortinet uses the malware statistics collected in this manner to improve the performance of the FortiGate services and to display statistics on the [Fortinet Support](#) website for customers registered FortiGate devices.

Fortinet may also publish or share statistics or results derived from this malware data with various audiences. The malware statistics shared in this way do not include any customer data.

To enable, disable, and/or customize how often statistics are sent to FortiGuard, use the following command:

CLI syntax

```
config system global
  set fds-statistics {enable | disable}
  set fds-statistics-period <minutes>
end
```

In addition to secure submission of statistics to FortiGuard, there are other mechanisms in place to prevent unauthorized FortiGuard updates from clients:

- The server certificate has to be authenticated by FortiGates, and it only trusts Fortinet's root certificate.
- Proprietary encryption (including FGCP, an application-level proprietary protocol) that only Fortinet's own servers/devices can prepare.

FortiGates can only accept data from Fortinet's own list of servers, although the list can be updated through previously connected servers. DNS is used on the initial server, but all other servers are provided by a list that is updated through SSL, meaning that only FortiGates accept data from those servers.

Automatic update at every GUI login

FortiGates running FortiOS 5.6.1 and above may perform automatic "update now" updates when one of the "core" licenses is unavailable: Application Control, IPS, or AntiVirus. Please note that this automatic update is triggered even if the following CLI command is set:

```
config system autoupdate schedule
  set status disable
end
```

CLI Syntax

The following section contains commands to control FortiGuard.

system autoupdate push-update

The following command will set the FDN push update port.

```
config system.autoupdate push-update
  set port <integer>
end
```

system autoupdate tunneling

The following command will set the proxy server port that the FortiGate will use to connect to the FortiGuard Distribution Network (FDN).

```
config system.autoupdate tunneling
    set port <integer>
end
```

system fortiguard

The following command will set the port by which scheduled FortiGuard service updates will be received.

```
config system fortiguard
    set port {53 | 8888 | 80}
end
```

webfilter fortiguard

The following command will close ports used for HTTPS/HTTP override authentication and disable user overrides:

```
config webfilter fortiguard
    set close-ports {enable | disable}
end
```

For more information, including FortiGuard execute commands used to manage FortiCloud domains and operations, see the [CLI Reference](#).

FortiLink

FortiGate units can be used to remotely manage FortiSwitch units, which is also known as using a FortiSwitch in FortiLink mode. FortiLink defines the management interface and the remote management protocol between the FortiGate and FortiSwitch.

Different FortiGate models support remote management for varying numbers of FortiSwitches, as shown below:

FortiGate	Number of FortiSwitches
Up to FortiGate 98 and FortiGate VM01	8
FortiGate 100 to 280 and FortiGate VM02	24
FortiGate 300 to 5xx	48
FortiGate 600 to 900 and FortiGate VM04	64
FortiGate 1000 and up	128
FortiGate-3000 and up, and FortiGate VM08 and up	300

Supported FortiSwitch models

The following table shows the FortiSwitch models that support FortiLink mode when paired with the corresponding FortiGate models and the listed minimum software releases.

FortiSwitch	FortiGate	Earliest FortiSwitchOS	Earliest FortiOS
FS-224D-POE	FGT-90D (WiFi/POE)	3.0.0	5.2.2
FS-108D-POE	FGT-60D (all)	3.0.1	5.2.3
FSR-112D-POE	FGR-90D	3.0.1	5.2.3
FS-124D	FGT-90D + FGT-60D	3.0.1	5.2.3
FS-124D-POE	FGT-90D + FGT-60D	3.0.1	5.2.3
FS-224D-FPOE	FGT-90D + FGT-60D	3.0.1	5.2.3

Note that **all** FortiSwitches above also support FortiLink mode when paired with the following FortiGate models: 100D, 140D (POE, T1), 200D, 240D, 280D (POE), 600C, 800C, and 1000C.

FortiLink ports for each FortiSwitch model

Each FortiSwitch model provides one designated port for the FortiLink connection. The table below lists the FortiLink port for each model:

FortiSwitch model	Port for FortiLink connection
FS-28C	WAN port 1
FS-324B-POE	Management Port
FS-448B (10G only)	WAN port (uplink 1)
FS-348B	Last port (port 48)
For all D-series switches, use the last (highest number) port for FortiLink. For example:	
FS-108D-POE	Last port (port 10)
FSR-112D-POE	Last port (port 12)
FS-124D	Last port (port 26). May require an SFP module.*
FS-224D-POE	Last port (port 24)
FS-224D-FPOE	Last port (port 28). May require an SFP module.*

* FortiSwitch 3.3.1 and later releases support the use of an RJ-45 port for FortiLink. For additional information, visit the [Fortinet Support](#) website.

FortiLink ports for each FortiGate model

The following table shows the ports for each model of FortiGate that can be FortiLink-dedicated.

FortiGate model	Port for FortiLink connection
FGT-90D, FGT-90D-POE, FWF-90D, FWF-90D-POE	port1 - port14
FGT-60D, FGT-60D-POE, FWF-60D, FWF-60D-POE	port1 - port7
FGT-100D	port1 - port16
FGT-140D, 140D-POE, 140D-POE-T1	port1 - port36
FGT-200D	port1 - port16
FGT-240D	port1 - port40
FGT-280D, FGT-280D-POE	port1 - port84
FGT-600C	port3 - port22

FortiGate model	Port for FortiLink connection
FGT-800C	port3 - port24
FGT-1000C	port3 - port14, port23, port24

Auto-discovery of the FortiSwitch ports

In releases FortiSwitchOS 3.3.0 and beyond, the D-series FortiSwitch models support FortiLink auto-discovery, which is automatic detection of the port connected to the FortiGate.

You can use any of the switch ports for FortiLink. Use the following FortiSwitch CLI commands to configure a port for FortiLink auto-discovery:

```
config switch interface
edit <port>
set auto-discovery-fortilink enable
end
```

Note that some FortiSwitch ports are enabled for auto-discovery by default.

Each FortiSwitch model provides a set of ports that are enabled for FortiLink auto-discovery by default. If you connect the FortiLink using one of these ports, no switch configuration is required.

In general (in FortiSwitchOS 3.4.0 and later releases), the last four ports are the default auto-discovery FortiLink ports. The table below lists the default auto-discovery ports for each switch model:

FortiSwitch model	Default Auto-FortiLink ports
FS-108D	ports 9 and 10
FSR-112D	ports 9, 10, 11, and 12
FS-124D, FS-124D-POE	ports 23, 24, 25, and 26
FS-224D-POE	ports 21, 22, 23, and 24
FS-224D-FPOE	ports 25, 26, 27, and 28
FS-248D-POE	ports 49, 50, 51, and 52
FS-248D-FPOE	ports 49, 50, 51, and 52
FS-424D, FS-424D-POE, FS-424D-FPOE	ports 25 and 26
FS-448D, FS-448D-POE, FS-448D-FPOE	ports 49, 50, 51, and 52
FS-524D, FS-524D-FPOE	ports 25, 26, 27, 28, 29, and 30
FS-548D, FS-548D-FPOE	ports 49, 50, 51, 52, 53, and 54
FS-1024D, FS-1048D, FS-3032D	all ports

You can also run the `show switch interface` CLI command on the FortiSwitch to see the ports that have auto-discovery enabled.

Adding a managed FortiSwitch to the FortiGate

The following steps show how to add a new managed FortiSwitch using the FortiGate GUI or the CLI.



For FortiSwitchOS releases prior to 3.3.0, you must [Set the FortiSwitch to remote management mode](#) before following the steps below.

Using the FortiGate GUI:

1. Connect a cable from the designated FortiSwitch port to an unused port on the FortiGate. Refer to [FortiLink ports for each FortiSwitch model](#) for additional information.
2. Go to **Network > Interfaces** and edit an internal port on the FortiGate.
3. Set **Addressing mode** to **Dedicated to FortiSwitch** and select **OK**.
4. As of FortiOS 5.4.0, the **Managed FortiSwitch** GUI option can only be accessed by enabling it through the CLI console.

Open the CLI console and enter the following command to make the switch controller available in the GUI, and to set the reserved subnetwork for the controller:

```
config system global
    set switch-controller enable
    set switch-controller-reserved-network 169.254.254.0 255.255.255.0
end
```

5. Go to **WiFi & Switch Controller > Managed FortiSwitch**. The new FortiSwitch should now be displayed in the table.
6. Right-click on the FortiSwitch and select **Authorize**.

Using the FortiGate CLI:

Note that, for the example shown below, the FortiGate's port1 is configured as the FortiLink port.

1. If required, remove port1 from the **lan** interface:

```
config system virtual-switch
    edit lan
        config port
            delete port1
        end
    end
end
```

2. Configure the interface for port1:

```
config system interface
    edit port1
        set ip 172.20.120.10 255.255.255.0
        set allowaccess capwap
        set vlanforward enable
    end
end
```

3. Configure an NTP server on port1:

```
config system ntp
```

```
set server-mode enable
set interface port1
end
```

4. Authorize the FortiSwitch unit as a managed switch (note that FortiSwitch will reboot once you issue the command below):

```
config switch-controller managed-switch
edit FS224D3W14000370
set fsw-wan1-admin enable
end
end
```

5. Configure a DHCP server on port1:

```
config system dhcp server
edit 0
set netmask 255.255.255.252
set interface port1
config ip-range
edit 0
set start-ip 169.254.254.2
set end-ip 169.254.254.50
end
set vci-match enable
set vci-string FortiSwitch
set ntp-service local
end
end
```

Set the FortiSwitch to remote management mode

Use the FortiSwitch GUI or the CLI to set the remote management mode.

Note that the following steps are not necessary for FortiSwitchOS releases 3.3.0 or later.

Using the FortiSwitch GUI:

1. Go to **Dashboard > Main > System Information**.
2. Beside **Operation Mode**, select **Change**.
3. Change **Management Mode** to **FortiGate Remote Management** and select **OK**.
4. A warning will appear asking if you wish to continue. Select **OK**.

Using the FortiSwitch CLI:

```
config system global
set switch-mgmt-mode fortilink
end
```

Configuring the FortiSwitch remote management port

If the FortiSwitch model has a dedicated management port, you can configure remote management to the FortiSwitch. In FortiLink mode, the FortiGate is the default gateway, so you need to configure an explicit route for the FortiSwitch management port.

To do this, from the FortiSwitch CLI, enter the following command:

```
config router static
edit 1
```



```
set device mgmt
set gateway <router_IP_address>
set dst <router_subnet> <subnet_mask>
end
end
```

Configuring FortiLink LAG

Starting with FortiOS 5.4.0 and FortiSwitchOS 3.3.0, you can configure the FortiLink as a Link Aggregation Group (LAG) to provide increased bandwidth between the FortiGate and FortiSwitch.

Connect any two ports on the FortiGate to two ports on the FortiSwitch. Make sure that you use the designated FortiLink port as one of the ports on the switch.

To configure the FortiLink as a LAG on the FortiGate, create a trunk (of type `fortilink`) with the two ports that you connected to the switch:

```
config system interface
  edit "fortilink"
    set vdom root
    set allowaccess ping capwap http https
    set type fortilink
    set member port4 port5
    set snmp-index 17
    set lacp-mode static
  next
end
config system ntp
  set ntpsync enable
  set syncinterval 60
  set server-mode enable
  set interface "fortilink"
end
```

There is no specific configuration required for the LAG on the switch.

FortiOS WAN optimization

Organizations with multiple locations or businesses using the cloud can provide license-free WAN optimization using FortiOS.

WAN Optimization is a comprehensive solution that maximizes your WAN performance and provides intelligent bandwidth management and unmatched consolidated security performance. WAN optimization reduces your network overhead and removes unnecessary traffic for a better overall performance experience. Efficient use of bandwidth and better application performance will remove the need for costly WAN link upgrades between data centers and other expensive solutions for your network traffic growth.

WAN optimization is available on FortiGate models with internal storage that also support SSL acceleration. Internal storage includes high-capacity internal hard disks, AMC hard disk modules, FortiGate Storage Modules (FSMs) or over 4 GB of internal flash storage.

WAN optimization tunnels use port 7810.



The following features below are available through WAN optimization:

Protocol optimization

Protocol optimization is effective for applications designed for the LAN that do not function well on low bandwidth, high latency networks. FortiOS protocol optimization improves the efficiency of CIFS, FTP, HTTP, MAPI, and general TCP sessions.

CIFS, for example, requires many background transactions to successfully transfer a single file. When transferring the file, CIFS sends small chunks of data and waits sequentially for each chunk's arrival and acknowledgment before sending the next chunk. This large amount of requests and acknowledgements of traffic can delay transfers. WAN Optimization removes this complexity and improves the efficiency of transferring the file.

TCP protocol optimization uses techniques such as SACK support, window scaling and window size adjustment, and connection pooling to remove common WAN TCP bottlenecks.

Regular bandwidth usage**Improved bandwidth usage with FortiGate protocol optimization****Byte caching**

Byte caching improves caching by accelerating the transfer of similar, but not identical content. Byte caching reduces the amount of data crossing the WAN when multiple different emails with the same or similar attachments or different versions of an attachment are downloaded from a corporate email server to different locations over the WAN.

Byte caching breaks large units of application data, such as email attachments or file downloads, into smaller chunks of data. Each chunk of data is labeled with a hash, and chunks with their respective hashes are stored in a database on the local FortiGate unit. When a remote user requests a file, WAN optimization sends the hashes, rather than the actual data. The FortiGate unit at the other end of the WAN tunnel reassembles the data from its own hash database, only downloading the chunks it is missing. Deduplication, or the process of eliminating duplicate data, will reduce space consumption.

Byte caching is not application specific, and assists by accelerating all protocols supported by WAN optimization.

Web caching

WAN optimization reduces download times of content from central files repositories through web caching. FortiOS Web caching stores remote files and web pages on local FortiGate devices for easy local access to commonly accessed files. There is little impact on the WAN, resulting in reduced latency for those requesting the files.

In addition, web caching also recognizes requests for Windows or MS Office updates, and downloads the new update file in the background. Once downloaded to the cache, the new update file is available to all users, and all subsequent requests for this update are rapidly downloaded from the cache.

Traffic shaping

Controls data flow for specific applications, giving administrators the flexibility to choose which applications take precedence over the WAN. A common use case of traffic shaping would be to prevent one protocol or application from flooding a link over other protocols deemed more important by the administrator.

SSL acceleration

SSL is used by many organizations to keep WAN communications private. WAN Optimization boosts SSL acceleration properties of FortiGate FortiASIC hardware by accelerating SSL traffic across the WAN. The FortiGate unit handles SSL encryption/decryption for corporate servers providing SSL encrypted connections over the WAN.

Explicit web proxy server

Allows users on the internal network to browse the Internet through the explicit web proxy server.

Explicit FTP proxy server

Allows users on the internal network to access FTP servers through the explicit FTP proxy server.

Reverse proxy

The web and FTP proxies can be configured to protect access to web or FTP servers that are behind the FortiGate using a reverse proxy configuration. Reverse proxies retrieve resources on behalf of a client from one or more servers. These resources are then returned to the client as if they originated from the proxy server.

WCCP

The Web Cache Communication Protocol (WCCP) allows you to offload web caching to redundant web caching servers. This traffic redirection helps to improve response time and optimize network resource usage.

WAN optimization and HA

You can configure WAN optimization on a FortiGate HA cluster. The recommended HA configuration for WAN optimization is active-passive mode. Also, when the cluster is operating, all WAN optimization sessions are processed by the primary unit only. Even if the cluster is operating in active-active mode, HA does not load-balance WAN optimization sessions. HA also does not support WAN optimization session failover.

Configuring an explicit proxy with WAN optimization web caching

For this configuration, all devices on the wireless network will be required to connect to the proxy at port 8080 before they can browse the Internet. WAN Optimization web caching is added to reduce the amount of Internet bandwidth used and improve web browsing performance.

Enabling WAN optimization and configuring the explicit web proxy for the wireless interface

1. Go to **System > Feature Visibility**. Ensure that **Explicit Proxy** is enabled.
To make WAN Optimization and Web Caching settings available in the GUI, enter the following CLI command:

```
config system settings
    set gui-wanopt-cache enable
end
```
1. Go to **Network > Interfaces**, edit the wireless interface and select **Enable Explicit Web Proxy**.
2. Go to **Network > Explicit Proxy**. Enable **Explicit Web Proxy**. Make sure that **Default Firewall Policy Action** is set to **Deny**.

Adding an explicit web proxy policy

1. Go to **Policy & Objects > Proxy Policy** and create a new policy.
2. Set **Proxy Type** to **Explicit Web**, the **Outgoing Interface** to the Internet-facing interface, and the remaining fields as required.

For more detailed information, see [General web proxy configuration steps](#).

Configuring devices on the wireless network to use the web proxy

To use the web proxy, all devices on the wireless network must be configured to use the explicit proxy server. The IP address of the server is the IP address of the FortiGate's wireless interface (for example, 10.10.80.1) and the port is 8080. Some browsers may have to be configured to use the device's proxy settings.

For Windows 10, right-click the Windows start-icon and select **Network Connections**. Select **Proxy** and configure the proxy settings.

For Windows Vista/7/8, open **Internet Properties**. Go to **Connections > LAN Settings** and enable and configure the **Proxy Server**.

For Mac OS X, open **System Preferences > Network > Wi-Fi > Advanced > Proxies**. Select **Web Proxy (HTTP)** and configure the proxy settings.

For iOS, go to **Settings > Wi-Fi**. Edit the wireless network. Scroll down to **HTTP PROXY**, select **Manual**, and configure the proxy settings.

For Android, in WiFi network connection settings, edit the wireless network. Select **Show advanced options**, configure a **Manual** proxy, and enter the proxy settings.

Force HTTP and HTTPS traffic to use the web proxy

Block HTTP and HTTPS access to the Internet from the wireless network so that the only path to the Internet is through the explicit proxy. You can edit or delete policies that allow HTTP or HTTPS access. You can also add a policy to the top of the list that **Denies** HTTP and HTTPS traffic.

FSSO - Fortinet Single Sign-On

Fortinet Single Sign-On (FSSO), formerly known as FortiGate Server Authentication Extension (FSAE), is the authentication protocol by which users can transparently authenticate to FortiGate, FortiAuthenticator, and FortiCache devices. The FortiAuthenticator unit identifies users based on their authentication from a different system, and can be authenticated via numerous methods:

- Users can authenticate through a web portal and a set of embeddable widgets.
- Users with FortiClient Endpoint Security installed can be automatically authenticated through the FortiClient SSO Mobility Agent.
- Users authenticating against Active Directory can be automatically authenticated.
- RADIUS Accounting packets can be used to trigger an FSSO authentication.
- Users can be identified through the FortiAuthenticator API. This is useful for integration with third-party systems.

Below are the TCP/UDP ports used by the multiple FSSO modes:

Purpose	Protocol/Port
LDAP group membership lookup (Global Catalog)	TCP/3268
LDAP domain controller discovery and group membership lookup	TCP/389
DC Agent keepalive and push logon info to CA	UDP/8002
CA keepalive and push logon info to FortiGate	TCP/8000
NTLM	TCP/8000
CA DNS	UDP/53
Workstation check, polling mode (preferred method)	TCP/445
Workstation check, polling mode (fallback method)	TCP/135, TCP/139, UDP/137
Remote access to logon events	TCP/445
Group lookup using LDAP	TCP/389
Group lookup using LDAP with global catalog	TCP/3268
Group lookup using LDAPS	TCP/636
Resolve FSSO server name	UDP/53

Configuring the FortiAuthenticator

The FortiAuthenticator unit can be integrated with external network authentication systems, such as RADIUS, LDAP, Windows AD, and FortiClients to poll user logon information and send it to the FortiGate unit.

To configure FortiAuthenticator polling:

1. Go to **Fortinet SSO Methods > SSO > General**.
2. In the **FortiGate** section, leave **Listening port** set to **8000**, unless your network requires you to change this. The FortiGate unit must allow traffic on this port to pass through the firewall. Optionally, you can set the **Login expiry** time (default is **480** minutes, or eight hours). This is the length of time users can remain logged in before the system logs them off automatically.
3. Select **Enable authentication** and enter the **Secret key**. Be sure to use the same secret key when configuring the FSSO Agent on FortiGate units.
4. In the **Fortinet Single Sign-On (FSSO)** section, enter the following information:

Enable Windows event log polling (e.g. domain controllers/Exchange servers)	Select for integration with Windows Active Directory
Enable RADIUS Accounting SSO clients	Select if you want to use a Remote RADIUS server.
Enable Syslog SSO	Select for integration with Syslog server.
Enable FortiClient SSO Mobility Agent Service	<p>Once enabled, also select Enable authentication to enable SSO by clients running FortiClient Endpoint Security.</p> <p>Enter the Secret key. Be sure to use the same secret key in the FortiClient Single Sign-On Mobility Agent settings.</p>

5. Select **OK**.

For more detailed information for each available setting, see the [FortiAuthenticator Administration Guide](#).

Configuring the FortiGate

The FortiAuthenticator unit needs to be added to the FortiGate as an SSO agent that will provide user logon information.

To add a FortiAuthenticator unit as SSO agent:

1. Go to **Security Fabric > Fabric Connectors** and select **Create New**.
2. Under **SSO/Identity**, select **Fortinet Single-Sign-On Agent**.
3. Enter a **Name**, set **Primary FSSO Agent** either to the IP address of the FortiAuthenticator unit or a name, and enter a **Password**.
4. Set **Collector Agent AD access mode** to either **Standard**, where you can specify **Users/Groups**, or **Advanced**, where you can specify an **LDAP Server**.
5. Select **OK**.

The FortiGate unit receives a list of user groups from the FortiAuthenticator unit or LDAP server. When you open the server, you can see the list of groups. You can use the groups in identity-based security policies.

FSSO user groups

You can only use FortiAuthenticator SSO user groups directly in identity-based security policies. You must create an FSSO user group, then add FortiAuthenticator SSO user groups to it. These FortiGate FSSO user groups will then become available for selection in identity-based security policies.

To create an FSSO user group:

1. Go to **User & Device > User Groups** and select **Create New**.
2. Enter a **Name** for the group.
3. Set **Type** to **Fortinet Single Sign-On (FSSO)**.
4. Add **Members**. The groups available to add as members are SSO groups provided by SSO agents.
5. Select **OK**.

Configuring the FortiClient SSO Mobility Agent

In order for the user to successfully set up the SSO Mobility Agent in FortiClient, they must know the FortiAuthenticator IP address and pre-shared key/secret.

To configure FortiClient SSO Mobility Agent:

1. In FortiClient, go to **File > Settings**.
2. Under **Advanced**, select **Enable Single Sign-On mobility agent**.
3. In **Server address**, enter the IP address of the FortiAuthenticator.
4. In **Customize port**, enter the listening port number specified on the FortiAuthenticator unit. You can omit the port number if it is **8005**.
5. Enter the **Pre-shared key**.
6. Select **OK**.

For more detailed FSSO configurations, including with Windows AD, Citrix, Novell eDirectory, and more, see the [Authentication](#) guide.

CLI Syntax

The following section contains commands to control FSSO.

user fsso

The following command will set the server address, port, and password for multiple FSSO agents.

```
config user fsso
  edit <name_str>
    set name <string>
    set [server | server2 | server3 | server4 | server5] <string>
    set [port | port2 | port3 | port4 | port5] <integer>
    set [password | password2 | password3 | password4 | password5] <password>
  end
```

user fsso-polling

The following command will set the Active Directory server port.

```
config user fsso-polling
```



```
edit <name_str>
    set port <integer>
end
```

OFTP - Optimized Fabric Transfer Protocol

The Optimized Fabric Transfer Protocol (OFTP) is used when information is synchronized between FortiAnalyzer and FortiGate. Remote logging and archiving can be configured on the FortiGate to send logs to a FortiAnalyzer (and/or FortiManager) unit.

OFTP listens on ports TCP/514 and UDP/514.

You can connect to a FortiAnalyzer unit from a FortiGate unit using Automatic Discovery, so long as both units are on the same network. Connecting these devices in this way does not use OFTP. Instead, the Fortinet Discovery Protocol (FDP) is used to locate the FortiAnalyzer unit.

When you select Automatic Discovery, the FortiGate unit uses HELLO packets to locate any FortiAnalyzer units that are available on the network within the same subnet. When the FortiGate unit discovers the FortiAnalyzer unit, the FortiGate unit automatically enables logging to the FortiAnalyzer unit and begins sending log data.

CLI command - To connect to FortiAnalyzer using automatic discovery:

```
config log fortianalyzer setting
  set status [enable | disable]
  set server <ip_address>
  set gui-display [enable | disable]
  set address-mode auto-discovery
end
```



If your FortiGate unit is in Transparent mode, the interface using the automatic discovery feature will not carry traffic.

To send logs from FortiGate to FortiAnalyzer:

1. Go to **Log & Report > Log Settings** and enable **Send logs to FortiAnalyzer/FortiManager** (under **Remote Logging and Archiving**).
2. Enter the FortiAnalyzer unit's IP address in the **IP address** field provided.
3. For **Upload option**, select either **Real Time** to upload logs as they come across the FortiGate unit, or **Every Minute**, or **Every 5 Minutes**.
4. Logs sent to FortiAnalyzer can be encrypted by enabling **SSL encrypt log transmission**.

FortiClient EMS - Enterprise Management Server

FortiClient Enterprise Management Server (FortiClient EMS) is a security management solution that enables scalable and centralized management of multiple endpoints (computers). FortiClient EMS provides efficient and effective administration of endpoints running FortiClient. It provides visibility across the network to securely share information and assign security profiles to endpoints. It is designed to maximize operational efficiency and includes automated capabilities for device management and troubleshooting.

FortiClient EMS is designed to meet the needs of small to large enterprises that deploy FortiClient on endpoints. Benefits of deploying FortiClient EMS include:

- Remotely deploying FortiClient software to Windows PCs.
- Updating profiles for endpoint users regardless of access location, such as administering antivirus, web filtering, VPN, and signature updates.
- Administering FortiClient endpoint registrations, such as accepting, deregistering, and blocking registrations.
- Managing endpoints, such as status, system, and signature information.
- Identifying outdated versions of FortiClient software.

Required services

You must ensure that required ports and services are enabled for use by FortiClient EMS and its associated applications on your server. The required ports and services enable FortiClient EMS to communicate with clients and servers running associated applications.

Communication	Service	Protocol	Port
FortiClient endpoint registration	File transfers	TCP	8013 (default)
Computer browser service	Enabled		
Samba (SMB) service <ul style="list-style-type: none">• During FortiClient deployment, endpoints may connect to the FortiClient EMS server using the SMB service.	Enabled		445
Distributed Computing Environment / Remote Procedure Calls (DCE- RPC) <ul style="list-style-type: none">• The FortiClient EMS server connects to the endpoints using RPC for FortiClient deployment.	Enabled		135
Active Directory server connection	When used as a default connection		389
Windows	HTTP	TCP	80

Communication	Service	Protocol	Port
Internet Information Services (IIS)	HTTPS	TCP	443, 10443
SQL server			

For more informationn about FortiClient EMS, including other requirements, installation, and management, see the [FortiClient EMS - Administration Guide](#).

Chapter 9 - FortiWiFi and FortiAP Configuration Guide

This FortiOS Handbook chapter contains the following sections:

[Introduction to wireless networking](#) explains the basic concepts of wireless networking and how to plan your wireless network.

[Configuring a WiFi LAN](#) explains how to set up a basic wireless network, prior to deploying access point hardware.

[Access point deployment](#) explains how to deploy access point hardware and add it to your wireless network configuration.

[Wireless mesh](#) explains how to configure a Wi-Fi network where access points are connected to the Wi-Fi controller wirelessly instead of by Ethernet.

[Combining WiFi and wired networks with a software switch](#) shows how to use the FortiAP Wi-Fi-Ethernet bridge feature.

[Protecting the WiFi network](#) explains the Wireless Intrusion Detection System (WIDS).

[Wireless network monitoring](#) explains how to monitor your wireless clients and how to monitor other wireless access points, potentially rogues, in your coverage area.

[Configuring wireless network clients](#) explains how to configure typical wireless clients to work with a WPA-Enterprise protected network.

[Wireless network examples](#) provides two examples. The first is a simple Wi-Fi network using automatic configuration. The second is a more complex example of a business with two Wi-Fi networks, one for employees and another for guests or customers.

[Using a FortiWiFi unit as a client](#) explains how to use a FortiWiFi unit as a wireless client to connect to other Wi-Fi networks. This connection can take the place of an Ethernet connection where wired access to a network or to the Internet is not available.

[Support for location-based services](#) explains how Fortinet supports location-based services that collect information about devices near FortiGate-managed access points, even if the devices don't associate with the network.

[Reference](#) provides information about Wi-Fi radio channels.

What's new in FortiOS 6.0.1

The following new wireless features added in FortiOS 6.0.1. Click on a link to navigate to that section for further information.

- ["Locate a FortiAP with LED blinking" on page 1138](#)

What's new in FortiOS 6.0

The following list contains new wireless features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- ["MAC-auth-bypass for the captive-portal SSID" on page 1056](#)
- ["Session timeout interval for captive portal" on page 1062](#)
- ["Enforcing UTM policies on a local bridge SSID for managed smart APs" on page 1087](#)
- ["Multiple FortiAP firmware upgrades at once" on page 1131](#)
- ["Manual quarantine of devices on FortiAP \(tunnel mode\)" on page 1136](#)
- ["Host quarantine per SSID" on page 1137](#)
- ["Locate a FortiAP with LED blinking" on page 1138](#)
- ["Wireless controller optimization for large deployment - AP image upgrade" on page 1138](#)
- ["Control message off-loading and aeroscout enhancement" on page 1139](#)
- ["Preventing local bridge traffic from reaching the LAN" on page 1143](#)
- ["FortiAP-S bridge mode security profiles" on page 1144](#)
- ["DHCP snooping and option 82 \(circuit -id\) options for wireless access points" on page 1144](#)
- ["WiFi Channel Utilization" on page 1151](#)
- ["Support for extension information for wtp, vap, and station" on page 1211](#)

Introduction to wireless networking

This chapter introduces some concepts you should understand before working with wireless networks, describes Fortinet's wireless equipment, and then describes the factors you need to consider in planning deployment of a wireless network.

Wireless concepts

Wireless networking is radio technology, subject to the same characteristics and limitations as the familiar audio and video radio communications. Various techniques are used to modulate the radio signal with a data stream.

Bands and channels

Depending on the wireless protocol selected, you have specific channels available to you, depending on what region of the world you are in.

- IEEE 802.11b and g protocols provide up to 14 channels in the 2.400-2.500 GHz Industrial, Scientific and Medical (ISM) band.
- IEEE 802.11a,n (5.150-5.250, 5.250-5.350, 5.725–5.875 GHz, up to 16 channels) in portions of Unlicensed National Information Infrastructure (U-NII) band

Note that the width of these channels exceeds the spacing between the channels. This means that there is some overlap, creating the possibility of interference from adjacent channels, although less severe than interference on the same channel. Truly non-overlapping operation requires the use of every fourth or fifth channel, for example ISM channels 1, 6 and 11.

The capabilities of your wireless clients is the deciding factor in your choice of wireless protocol. If your clients support it, 5GHz protocols have some advantages. The 5GHz band is less used than 2.4GHz and its shorter wavelengths have a shorter range and penetrate obstacles less. All of these factors mean less interference from other access points, including your own.

When configuring your WAP, be sure to correctly select the Geography setting to ensure that you have access only to the channels permitted for WiFi use in your part of the world.

For detailed information about the channel assignments for wireless networks for each supported wireless protocol, see [Reference on page 1212](#).

Power

Wireless LANs operate on frequencies that require no license but are limited by regulations to low power. As with other unlicensed radio operations, the regulations provide no protection against interference from other users who are in compliance with the regulations.

Power is often quoted in dBm. This is the power level in decibels compared to one milliwatt. 0dBm is one milliwatt, 10dBm is 10 milliwatts, 27dBm, the maximum power on Fortinet FortiAP equipment, is 500 milliwatts. The FortiGate unit limits the actual power available to the maximum permitted in your region as selected by the WiFi controller country setting.

Received signal strength is almost always quoted in dBm because the received power is very small. The numbers are negative because they are less than the one milliwatt reference. A received signal strength of -60dBm is one millionth of a milliwatt or one nanowatt.

Antennas

Transmitted signal strength is a function of transmitter power and antenna gain. Directional antennas concentrate the signal in one direction, providing a stronger signal in that direction than would an omnidirectional antenna.

FortiWiFi units have detachable antennas. However, these units receive regulatory approvals based on the supplied antenna. Changing the antenna might cause your unit to violate radio regulations.

Security

There are several security issues to consider when setting up a wireless network.

Whether to broadcast SSID

It is highly recommended to broadcast the SSID. This makes connection to a wireless network easier because most wireless client applications present the user with a list of network SSIDs currently being received. This is desirable for a public network.

Attempting to obscure the presence of a wireless network by not broadcasting the SSID does not improve network security. The network is still detectable with wireless network “sniffer” software. Clients search for SSIDs that they know, leaking the SSID. Refer to [RFC 3370](#). Also, many of the latest Broadcom drivers do not support hidden SSID for WPA2.

Encryption

Wireless networking supports the following security modes for protecting wireless communication, listed in order of increasing security.

None — Open system. Any wireless user can connect to the wireless network.

WEP64 — 64-bit Web Equivalent Privacy (WEP). This encryption requires a key containing 10 hexadecimal digits.

WEP128 — 128-bit WEP. This encryption requires a key containing 26 hexadecimal digits.

WPA — 256-bit WiFi Protected Access (WPA) security. This encryption can use either the TKIP or AES encryption algorithm and requires a key of either 64 hexadecimal digits or a text phrase of 8 to 63 characters. It is also possible to use a RADIUS server to store a separate key for each user.

WPA2 — WPA with security improvements fully meeting the requirements of the IEEE 802.11i standard. Configuration requirements are the same as for WPA.

For best security, use the WPA2 with AES encryption and a RADIUS server to verify individual credentials for each user. WEP, while better than no security at all, is an older algorithm that is easily compromised. With either WEP or WAP, changing encryption passphrases on a regular basis further enhances security.

Separate access for employees and guests

Wireless access for guests or customers should be separate from wireless access for your employees. Each of the two networks can have its own SSID, security settings, firewall policies, and user authentication. This does not

require additional hardware. Both FortiWiFi units and FortiAP units support multiple wireless LANs on the same access point.

A good practice is to broadcast the SSID for the guest network to make it easily visible to users, but not to broadcast the SSID for the employee network.

Two separate wireless networks are possible because multiple virtual APs can be associated with an AP profile. The same physical APs can provide two or more virtual WLANs.

Captive portal

As part of authenticating your users, you might want them to view a web page containing your acceptable use policy or other information. This is called a captive portal. No matter what URL the user initially requested, the portal page is returned. Only after authenticating and agreeing to usage terms can the user access other web resources.

For more information about captive portals, see the Captive portals chapter of the FortiOS Authentication Guide.

Power

Reducing power reduces unwanted coverage and potential interference to other WLANs. Areas of unwanted coverage are a potential security risk. There are people who look for wireless networks and attempt to access them. If your office WLAN is receivable out on the public street, you have created an opportunity for this sort of activity.

Monitoring for rogue APs

It is likely that there are APs available in your location that are not part of your network. Most of these APs belong to neighboring businesses or homes. They may cause some interference, but they are not a security threat. There is a risk that people in your organization could connect unsecured WiFi-equipped devices to your wired network, inadvertently providing access to unauthorized parties. The optional On-Wire Rogue AP Detection Technique compares MAC addresses in the traffic of suspected rogues with the MAC addresses on your network. If wireless traffic to non-Fortinet APs is also seen on the wired network, the AP is a rogue, not an unrelated AP.

Decisions about which APs are rogues are made manually on the Rogue AP monitor page. For detailed information, see [Wireless network monitoring on page 1145](#).

Suppressing rogue APs

When you have declared an AP to be a rogue, you have the option of suppressing it. To suppress an AP, the FortiGate WiFi controller sends reset packets to the rogue AP. Also, the MAC address of the rogue AP is blocked in the firewall policy. You select the suppression action on the Rogue AP monitor page. For more information, see [Wireless network monitoring on page 1145](#).



Rogue suppression is available only when there is a radio dedicated to scanning. It will not function during background scanning for spectrum analysis.

Wireless Intrusion Detection (WIDS)

You can create a WIDS profile to enable several types of intrusion detection:

- Unauthorized Device Detection
- Rogue/Interfering AP Detection
- Ad-hoc Network Detection and Containment
- Wireless Bridge Detection
- Misconfigured AP Detection
- Weak WEP Detection
- Multi Tenancy Protection
- MAC OUI Checking

For more information, see [Protecting the WiFi network on page 1140](#).

Authentication

Wireless networks usually require authenticated access. FortiOS authentication methods apply to wireless networks the same as they do to wired networks because authentication is applied in the firewall policy.

The types of authentication that you might consider include:

- user accounts stored on the FortiGate
- user accounts managed and verified on an external RADIUS, LDAP or TACACS+ server
- Windows Active Directory authentication, in which users logged on to a Windows network are transparently authenticated to use the wireless network.

This FortiWiFi and FortiAP Configuration Guide provides some information about each type of authentication, but more detailed information is available in the Authentication chapter of the FortiOS Handbook.

What all of these types of authentication have in common is the definition of user groups to specify who is authorized. For each wireless LAN, you will create a user group and add to it the users who can use the WLAN. In the identity-based firewall policies that you create for your wireless LAN, you will specify this user group.

Some access points, including FortiWiFi units, support MAC address filtering. You should not rely on this alone for authentication. MAC addresses can be “sniffed” from wireless traffic and used to impersonate legitimate clients.

Wireless networking equipment

Fortinet produces two types of wireless networking equipment:

- [FortiWiFi units](#), which are FortiGate units with a built-in wireless access point/client
- [FortiAP units](#), which are wireless access points that you can control from any FortiGate unit that supports the WiFi Controller feature.

FortiWiFi units

A FortiWiFi unit can:

- Provide an access point for clients with wireless network cards. This is called Access Point mode, which is the default mode.
- or**
- Connect the FortiWiFi unit to another wireless network. This is called Client mode. A FortiWiFi unit operating in client mode can only have one wireless interface.

or

- Monitor access points within radio range. This is called Monitoring mode. You can designate the detected access points as Accepted or Rogue for tracking purposes. No access point or client operation is possible in this mode. But, you can enable monitoring as a background activity while the unit is in Access Point mode.

The Products section of the Fortinet web site (www.fortinet.com) provides detailed information about the FortiWiFi models that are currently available.

FortiAP units

FortiAP units are thin wireless access points are controlled by either a FortiGate unit or FortiCloud service.

FortiAP is a family of Indoor, Outdoor and Remote Access Point models supporting the latest single, dual, and triple stream MIMO 802.11ac and 802.11n technology, as well as 802.11g and 802.11a.

For large deployments, some FortiAP models support a mesh mode of operation in which control and data backhaul traffic between APs and the controller are carried on a dedicated WiFi network. Users can roam seamlessly from one AP to another.

In dual-radio models, each radio can function as an AP or as a dedicated monitor. The monitoring function is also available during AP operation, subject to traffic levels.

The [Products](#) section of the Fortinet web site (www.fortinet.com) provides detailed information about the FortiAP models that are currently available.

Automatic Radio Resource Provisioning

To prevent interference between APs, the FortiOS WiFi Controller includes the Distributed Automatic Radio Resource Provisioning (DARRP) feature. Through DARRP, each FortiAP unit autonomously and periodically determines the channel that is best suited for wireless communications. FortiAP units to select their channel so that they do not interfere with each other in large-scale deployments where multiple access points have overlapping radio ranges.

To enable ARRP - GUI

1. Go to **WiFi Controller > FortiAP Profiles** and edit the profile for your device.
2. In the Radio sections (Radio 1, Radio 2, etc.), enable **Radio Resource Provision**.
3. Click **OK**.

To enable ARRP - CLI

In this example, ARRP is enabled for both radios in the FAP321C-default profile:

```
config wireless-controller wtp-profile
  edit FAP321C-default
    config radio-1
      set darrp enable
    end
    config radio-2
      set darrp enable
    end
  end
end
```

Setting ARP timing

By default, ARP optimization occurs at a fixed interval of 1800 seconds (30 minutes). You can change this interval in the CLI. For example, to change the interval to 3600 seconds enter:

```
config wireless-controller timers
  set darrp-optimize 3600
end
```

Optionally, you can schedule optimization for fixed times. This enables you to confine ARP activity to a low-traffic period. Setting `darrp-optimize` to 0, makes `darrp-day` and `darrp-time` available. For example, here's how to set DARRP optimization for 3:00am every day:

```
config wireless-controller timers
  set darrp-optimize 0
  set darrp-day sunday monday tuesday wednesday thursday friday saturday
  set darrp-time 03:00
end
```

Both `darrp-day` and `darrp-time` can accept multiple entries.

Captive portals

A captive portal is a convenient way to authenticate web users on wired or WiFi networks.

This section describes:

- [Introduction to captive portals](#)
- [Configuring a captive portal](#)
- [Customizing captive portal pages](#)
- [Configuration example - captive portal WiFi access control](#)

Introduction to captive portals

You can authenticate your users on a web page that requests the user's name and password. Until the user authenticates successfully, the authentication page is returned in response to any HTTP request. This is called a captive portal.

After successful authentication, the user accesses the requested URL and can access other web resources, as permitted by security policies. Optionally, the captive portal itself can allow web access to only the members of specified user group.

The captive portal can be hosted on the FortiGate unit or on an external authentication server. You can configure captive portal authentication on any network interface, including WiFi and VLAN interfaces.

When a captive portal is configured on a WiFi interface, the access point initially appears open. The wireless client can connect to the access point with no security credentials, but sees only the captive portal authentication page.

WiFi captive portal types:

- **Authentication** — until the user enters valid credentials, no communication beyond the AP is permitted.
- **Disclaimer + Authentication** — immediately after successful authentication, the portal presents the disclaimer page—an acceptable use policy or other legal statement—to which the user must agree before proceeding.
- **Disclaimer Only** — the portal presents the disclaimer page—an acceptable use policy or other legal statement—to which the user must agree before proceeding. The authentication page is not presented.
- **Email Collection** — the portal presents a page requesting the user's email address, for the purpose of contacting the person in future. This is often used by businesses who provide free WiFi access to their customers. The authentication page is not presented.
- **MAC Bypass** — when clients are authenticated against their bridged SSID and their MAC addresses are known, they are redirected to the external captive portal.

Configuring a captive portal

Captive portals are configured on network interfaces. A WiFi interface does not exist until the WiFi SSID is created. You can configure a WiFi captive portal at the time that you create the SSID. Afterwards, the captive portal settings will also be available by editing the WiFi network interface in **System > Network > Interfaces**. On a physical (wired) network interface, you edit the interface configuration in **System > Network > Interfaces** and set **Security Mode** to **Captive Portal**.

To configure a WiFi captive portal - web-based manager:

1. Go to **WiFi & Switch Controller > SSID** and create your SSID.
If the SSID already exists, you can edit the SSID or you can edit the WiFi interface in **Network > Interfaces**.
2. Under **WiFi Settings**, for **Security Mode**, select **Captive Portal**.

WiFi Settings

SSID	<input type="text" value="fortinet"/>
Security Mode	<input type="text" value="Captive Portal"/>
Client Limit	<input type="checkbox"/>
Portal Type	<input checked="" type="radio"/> Authentication <input type="radio"/> Disclaimer + Authentication <input type="radio"/> Disclaimer Only <input type="radio"/> Email Collection
Authentication Portal	<input type="text" value="Local"/> <input checked="" type="text" value="External"/>
	<input type="text" value="example.com/captive/"/>
User Groups	<input type="text" value="+"/> +
Exempt Sources	<input type="text" value="+"/> +
Exempt Destinations/Services	<input type="text" value="+"/> +
Redirect after Captive Portal	<input checked="" type="radio"/> Original Request <input type="radio"/> Specific URL

3. Enter the following:

Portal Type	The portal can provide authentication and/or disclaimer, or perform user email address collection. See Introduction to captive portals on page 1054 .
Authentication Portal	Local - portal hosted on the FortiGate unit. Remote - enter FQDN or IP address of external portal.
User Groups	Select permitted user groups.
Exempt Sources	Select exempt lists whose members will not be subject to captive portal authentication.
Exempt Destinations/Services	
Redirect after Captive Portal	Select whether to have authenticated users navigate to their originally requested URL or be redirected to another/specific URL.

4. Select **OK**.

To configure an SSID with external-web enabled - CLI:

```
config wireless-controller vap
  edit "web-ext"
    set vdom "root"
    set ssid "web-ext"
    set security captive-portal
    set selected-usergroups "qnap"
    set security-exempt-list "wifi"
    set security-redirect-url " http://www.fortinet.com"
    set intra-vap-privacy enable
```

```
    set local-switching disable
    set external-web "192.168.234.51/portal.php"
  next
end
```

Note that the `external-web` entry is the URL of the external authentication web server. When this entry is not set, the FortiGate will use the local web server hosting the local login/splash page.

The external web URL is not explicitly set with HTTP/HTTPS - FortiGate uses the `auth-secure-http` entry under `config user` setting.

Exemption from the captive portal

A captive portal requires all users on the interface to authenticate. But some devices are not able to authenticate. You can create an exemption list of these devices. For example, a printer might need to access the Internet for firmware upgrades. Using the CLI, you can create an exemption list to exempt all printers from authentication.

```
config user security-exempt-list
  edit r_exempt
    config rule
      edit <id>
        set devices printer
      end
    end
  end
```

Furthermore, a walled garden firewall policy can be created:

```
config firewall policy
  edit <id>
    set captive-portal-exempt enable
    ...
  next
end
```

MAC Bypass for captive portal

It is possible to provide a MAC address bypass for authenticated clients. When clients are authenticated with bridged SSID and their MAC addresses are known, they are redirected to the External Captive Portal.

A new portal type has been added, under `config wireless-controller vap`, to provide successful MAC authentication Captive Portal functionality.

Syntax

```
config wireless-controller vap
  edit {name}
    set portal-type {cmcc-macauth}
  next
end
```

MAC-auth-bypass for the captive-portal SSID

Captive-portal SSID supports MAC-auth-bypass. If a client's MAC can be authenticated from local-user or RADIUS, then the client can bypass firewall authentication directly.

```
config wireless-controller vap
  edit <name>
    set security captive-portal
    set MAC-auth-bypass {enable | disable}
```

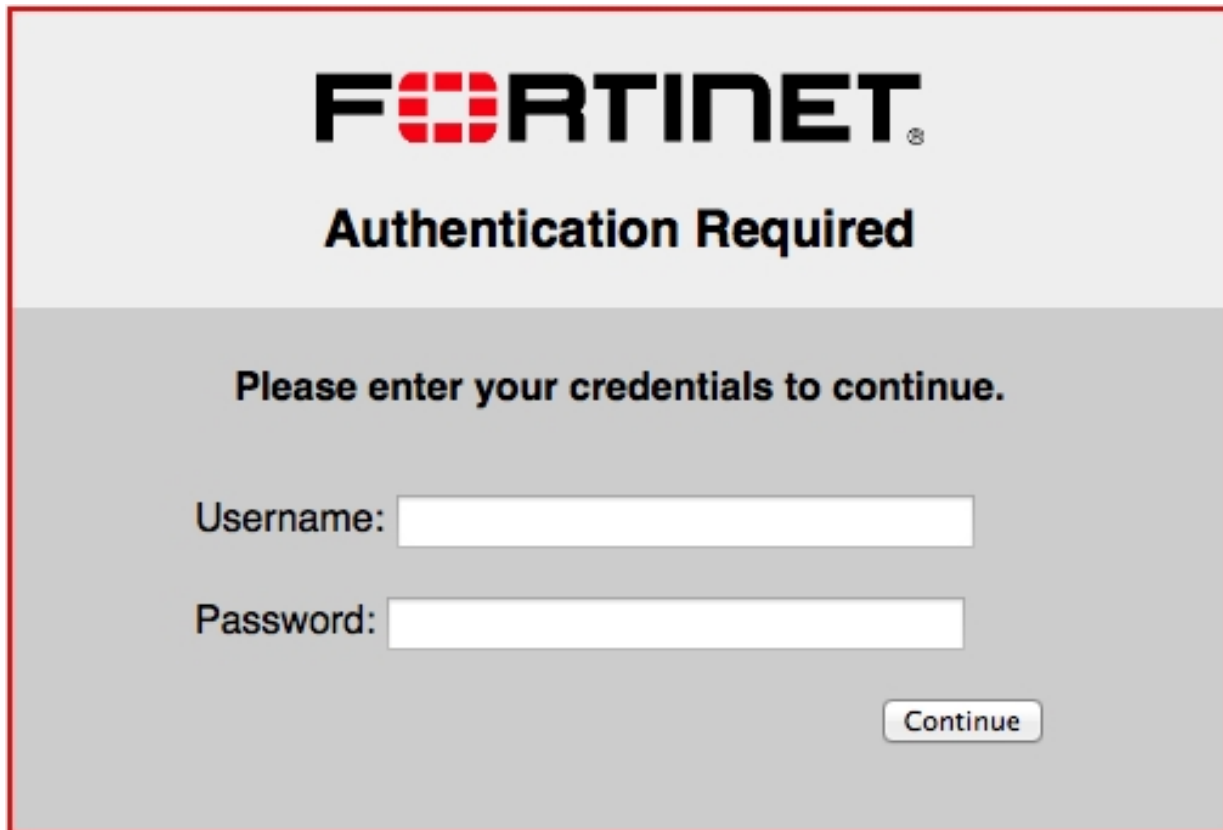
```
next
end
```

Customizing captive portal pages

These pages are defined in replacement messages. Defaults are provided. In the web-based manager, you can modify the default messages in the SSID configuration by selecting **Customize Portal Messages**. Each SSID can have its own unique portal content.

The captive portal contains the following default web pages:

- **Login page**—requests user credentials

The image shows a captive portal login page for Fortinet. At the top, the Fortinet logo is displayed in black and red. Below the logo, the text "Authentication Required" is centered in a bold, black font. Underneath this, a message "Please enter your credentials to continue." is also centered. There are two input fields: "Username:" followed by a white text box, and "Password:" followed by a white text box. At the bottom right of the form, there is a button labeled "Continue". The entire form is enclosed in a red border.

Typical modifications for this page would be to change the logo and modify some of the text.

You can change any text that is not part of the HTML code nor a special tag enclosed in double percent (%) characters.

There is an exception to this rule. The line "Please enter your credentials to continue" is provided by the `%%QUESTION%%` tag. You can replace this tag with text of your choice. Except for this item, you should not remove any tags because they may carry information that the FortiGate unit needs.

- **Login failed page**—reports that the entered credentials were incorrect and enables the user to try again.

The image shows a captive portal page for Fortinet. At the top, the Fortinet logo is displayed in black and red. Below the logo, the text "Authentication Failed" is centered in a bold, black font. Underneath this, a message reads "Firewall Authentication failed. Please try again." in a bold, black font. Below the message, there are two input fields: "Username:" followed by a white text box, and "Password:" followed by a white text box. At the bottom right of the form, there is a button labeled "Continue". The entire form is enclosed in a red border.

The Login failed page is similar to the Login page. It even contains the same login form. You can change any text that is not part of the HTML code nor a special tag enclosed in double percent (%) characters.

There is an exception to this rule. The line "Firewall authentication failed. Please try again." is provided by the %%FAILED_MESSAGE%% tag. You can replace this tag with text of your choice. Except for this item, you should not remove any tags because they may carry information that the FortiGate unit needs.

- **Disclaimer page**—is a statement of the legal responsibilities of the user and the host organization to which the user must agree before proceeding. (WiFi or SSL VPN only)

Terms and Disclaimer Agreement

You are about to access Internet content that is not under the control of the network access provider. The network access provider is therefore not responsible for any of these sites, their content or their privacy policies. The network access provider and its staff do not endorse nor make any representations about these sites, or any information, software or other products or materials found there, or any results that may be obtained from using them. If you decide to access any Internet content, you do this entirely at your own risk and you are responsible for ensuring that any accessed material does not infringe the laws governing, but not exhaustively covering, copyright, trademarks, pornography, or any other material which is slanderous, defamatory or might cause offence in any other way.

Do you agree to the above terms?

Yes, I agree

No, I decline

- **Declined disclaimer page**—is displayed if the user does not agree to the statement on the Disclaimer page. Access is denied until the user agrees to the disclaimer.



Changing images in portal messages

You can replace the default Fortinet logo with your organization's logo. First, import the logo file into the FortiGate unit and then modify the Login page code to reference your file.

To import a logo file:

1. Go to **System > Replacement Messages** and select **Manage Images**.
2. Select **Create New**.
3. Enter a **Name** for the logo and select the appropriate **Content Type**.
The file must not exceed 24 Kilo bytes.
4. Select **Browse**, find your logo file and then select **Open**.
5. Select **OK**.

To specify the new logo in the replacement message:

1. Go to **Network > Interfaces** and edit the interface.
The **Security Mode** must be **Captive Portal**.
2. Select the portal message to edit.
 - In SSL VPN or WiFi interfaces, in **Customize Portal Messages** click the link to the portal messages that you want to edit.
 - In other interfaces, make sure that **Customize Portal Messages** is selected, select the adjacent **Edit** icon, then select the message that you want to edit.
3. In the HTML message text, find the %%IMAGE tag.
By default it specifies the Fortinet logo: %%IMAGE:logo_fw_auth%%
4. Change the image name to the one you provided for your logo.
The tag should now read, for example, %%IMAGE:mylogo%%
5. Select **Save**.
6. Select **OK**.

Modifying text in portal messages

Generally, you can change any text that is not part of the HTML code nor a special tag enclosed in double percent (%) characters. You should not remove any tags because they may carry information that the FortiGate unit needs. See the preceding section for any exceptions to this rule for particular pages.

To modify portal page text

1. Go to **System > Network > Interfaces** and edit the interface.
The SSID **Security Mode** must be **Captive Portal**.
2. Select the portal message to edit.
 - In SSL VPN or WiFi interfaces, in **Customize Portal Messages** click the link to the portal messages that you want to edit.
 - In other interfaces, make sure that **Customize Portal Messages** is selected, select the adjacent **Edit** icon, then select the message that you want to edit.
3. Edit the HTML message text, then select **Save**.
4. Select **OK**.

Configuring disclaimer page for ethernet interface captive portals

While you can customize a disclaimer page for captive portals that connect via WiFi, the same can be done for wired connections. However, this can only be configured on the CLI Console, and only without configuring user groups.

When configuring a captive portal through the CLI, you may set `security-groups` to a specific user group. The result of this configuration will show an authentication form to users who wish to log in to the captive portal—**not** a disclaimer page. If you do not set any `security-groups` in your configuration, an "Allow all" status will be in effect, and the disclaimer page will be displayed for users.

The example CLI configuration below shows setting up a captive portal interface without setting security-groups, resulting in a disclaimer page for users:

```
config system interface
  edit "port1"
    set vdom "root"
    set ip 172.16.101.1 255.255.255.0
    set allowaccess ping https ssh snmp http
    set type physical
    set explicit-web-proxy enable
    set alias "LAN"
    set security-mode captive-portal
    set snmp-index 1
  next
end
```

Roaming support

Client devices can maintain captive portal authentication as they roam across different APs. By maintaining a consistent authentication, uninterrupted access to latency sensitive applications such as VoIP is ensured.

The Cloud will push a random per-AP Network encryption key to the AP. The key is 32 bytes in length, and is used in captive portal fast roaming. All APs of an AP Network will use the same encryption key. This key is randomly generated, and will be updated daily.

Session timeout interval for captive portal

The following syntax can be set to configure a session timeout interval in seconds for Captive Portal users. Set the range between 0 - 864000 (or no timeout to ten days). The default is set to 0.



This command is only available when **local-standalone** is set to **enable**, **security** is set to **captive-portal**, and then **portal-type** is set to either **cmcc** or **cmcc-macauth**.

Syntax

```
config wireless-controller vap
  edit <name>
    ...
    set captive-portal-session-timeout-interval <seconds>
  next
end
```

Configuration example - captive portal WiFi access control

In this scenario, you will configure the FortiGate for captive portal access so users can log on to your WiFi network.

You will create a user account (*rgreen*), add it to a user group (*employees*), create a captive portal SSID (*example-staff*), and configure a FortiAP unit. When the user attempts to browse the Internet, they will be redirected to the captive portal login page and asked to enter their username and password.

1. Enabling HTTPS authentication

Go to **User & Device > Authentication Settings**.

Under **Protocol Support**, enable **Redirect HTTP Challenge to a Secure Channel (HTTPS)**. This will make sure that user credentials are communicated securely through the captive portal.

2. Creating the user

Go to **User & Device > User Definition** and create a Local user (*rgreen*).

Create additional users if needed, and assign any authentication methods.

3. Creating the user group

Go to **User & Device > User Groups** and create a user group (*employees*).

Add **rgreen** to the group.

4. Creating the SSID

Go to **WiFi & Switch Controller > SSID** and configure the wireless network.
Some FortiGate models may show the GUI path as **WiFi & Switch Controller**.

Enter an **Interface Name** (*example-wifi*) and **IP/Network Mask**.

An address range under **DHCP Server** will be automatically configured.

Under **WiFi Settings**, enter an **SSID** name (*example-staff*), set **Security Mode** to **Captive Portal**, and add the **employees** user group.

5. Creating the security policy

Go to **Policy & Objects > Addresses** and create a new address for the SSID (*example-wifi-net*).

Set **Subnet/IP Range** to the same range set on the DHCP server in the previous step.

Set **Interface** to the SSID interface.

Go to **Policy & Objects > IPv4 Policy** and create a new policy for WiFi users to connect to the Internet.

Add both the **example-wifi-net** address and **employees** user group to **Source**.

6. Connecting and authorizing the FortiAP

Go to **Network > Interfaces** and edit an available interface.

Under **Address**, set **Addressing mode** to **Dedicated to Extension Device** and assign it an IP address.

Connect the FortiAP unit to the configured interface, then go to **WiFi & Switch Controller > Managed FortiAPs**.

The FortiAP is listed, but its **State** shows a greyed-out question mark — this is because it is waiting for authorization.

Highlight the FortiAP and select **Authorize**.

The question mark is now replaced by a red down-arrow — this is because it is authorized, but still offline.

Go to **WiFi & Switch Controller > FortiAP Profiles** and edit the profile.

For each radio, enable **Radio Resource Provision** and select your SSID.

Go back to **WiFi & Switch Controller > Managed FortiAPs** to verify that the FortiAP unit is online.

7. Results

When a user attempts to connect to the wireless network, they will be redirected to the captive portal login screen.

Members of the **employees** group must enter their **Username** and **Password**. The user will then be redirected to the URL originally requested.

On the FortiGate, go to **Monitor > WiFi Client Monitor** to verify that the user is authenticated.

Configuring a WiFi LAN

When working with a FortiGate WiFi controller, you can configure your wireless network before you install any access points. If you are working with a standalone FortiWiFi unit, the access point hardware is already present but the configuration is quite similar. Both are covered in this section.



On FortiGate model 30D, web-based manager configuration of the WiFi controller is disabled by default. To enable it, enter the following CLI commands:

```
config system global
    set gui-wireless-controller enable
end
```

The WiFi Controller and Switch Controller are enabled through the Feature Store (under **System > Feature Select**). However, they are separately enabled and configured to display in the GUI via the CLI.

To enable both WiFi and Switch controllers, enter the following:



```
config system global
    set wireless-controller enable
    set switch-controller enable
end
```

To enable the GUI display for both controllers, have also been separated:

```
config system settings
    set gui-wireless-controller enable
    set gui-switch-controller enable
end
```

If you want to connect and authorize external APs, such as FortiAP units, see the next chapter, [Access point deployment](#).

Overview of WiFi controller configuration

The FortiGate WiFi controller configuration is composed of three types of object, the SSID, the AP Profile and the physical Access Point.

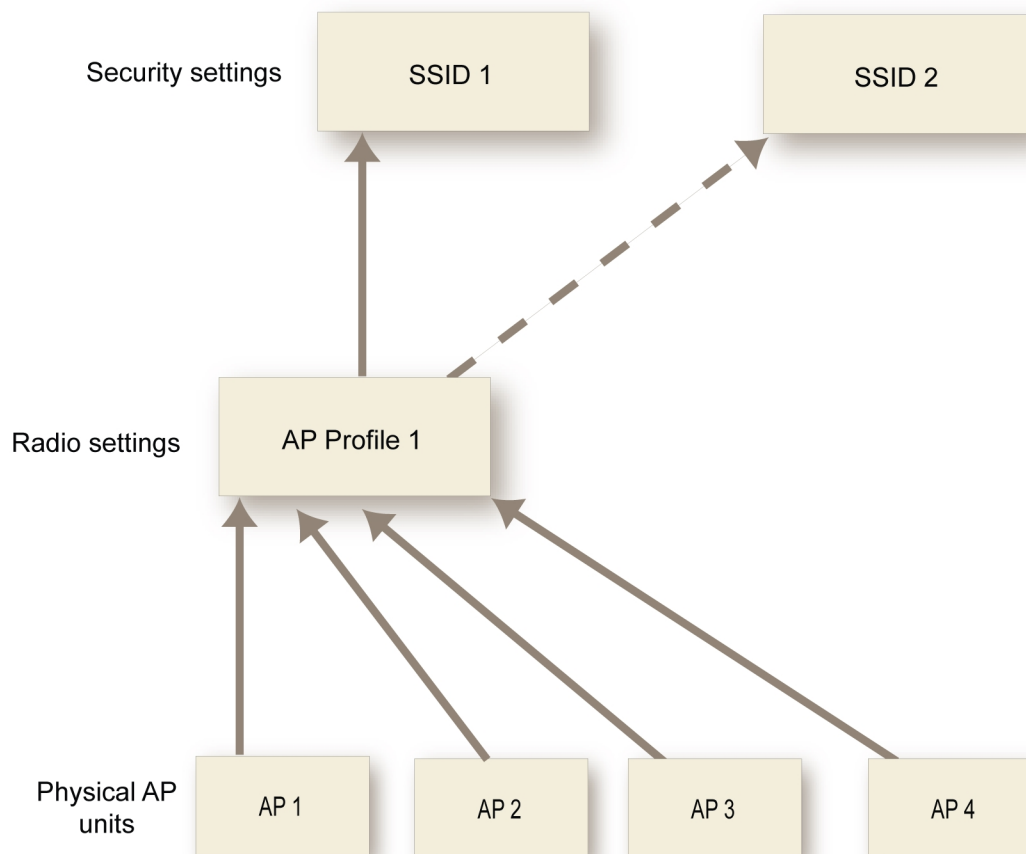
- An **SSID** defines a virtual wireless network interface, including security settings. One SSID is sufficient for a wireless network, regardless how many physical access points are provided. You might, however, want to create multiple SSIDs to provide different services or privileges to different groups of users. Each SSID has separate firewall policies and authentication. Each radio in an access point can support up to 8 SSIDs.

A more common use of the term SSID is for the identifier that clients must use to connect to the wireless network. Each SSID (wireless interface) that you configure will have an SSID field for this identifier. In Managed Access

Point configurations you choose wireless networks by SSID values. In firewall policies you choose wireless interfaces by their SSID name.

- An **AP Profile** defines the radio settings, such as band (802.11g for example) and channel selection. The AP Profile names the SSIDs to which it applies. Managed APs can use automatic profile settings or you can create AP profiles.
- **Managed Access Points** represent local wireless APs on FortiWiFi units and FortiAP units that the FortiGate unit has discovered. There is one managed access point definition for each AP device. An access point definition can use automatic AP profile settings or select a FortiAP Profile. When automatic profile settings are used, the managed AP definition also selects the SSIDs to be carried on the AP.

Conceptual view of FortiGate WiFi controller configuration



About SSIDs on FortiWiFi units

FortiWiFi units have a default SSID (wireless interface) named **wlan**. You can modify or delete this SSID as needed. As with external APs, the built-in wireless AP can be configured to carry any SSID.

The AP settings for the built-in wireless access point are located at **WiFi Controller > Local WiFi Radio**. The available operational settings are the same as those for external access points which are configured at **WiFi Controller > Managed FortiAPs**.

Process to create a wireless network

To set up your wireless network, you will need to perform the following steps:

- Make sure the FortiGate wireless controller is configured for your geographic location. This ensures that the available radio channels and radio power are in compliance with the regulations in your region.
- Optionally, if you don't want to use automatic AP profile settings, configure a FortiAP profile, specifying the radio settings and the SSIDs to which they apply.
- Configure one or more SSIDs for your wireless network. The SSID configuration includes DHCP and DNS settings.
- Configure the user group and users for authentication on the WLAN.
- Configure the firewall policy for the WLAN.
- Optionally, customize the captive portal.
- Configure access points.

Configuration of the built-in AP on FortiWiFi units is described in this chapter. Connection and configuration of FortiAP units is described in the next chapter, see [Access point deployment on page 1089](#).

Setting your geographic location

The maximum allowed transmitter power and permitted radio channels for WiFi networks depend on the region in which the network is located. By default, the WiFi controller is configured for the United States. If you are located in any other region, you need to set your location before you begin configuring wireless networks.

To change the location setting - CLI

To change the country to France, for example, enter

```
config wireless-controller setting
  set country FR
end
```

To see the list of country codes, enter a question mark ('?') instead of a country code.



Before changing the country setting, you must remove all FortiAP Profiles. To do this, go to **WiFi & Switch Controller > FortiAP Profiles**.

View all country and regcodes/regulatory domains

The following CLI command can be entered to view a list of the country and regcodes/regulatory Domains supported by Fortinet:

```
cw_diag -c all-countries
```

Below is a table showing a sample of the list displayed by entering this command:

Country-code	Region-code	Domain	ISO-name	Name
0	A	FCC3 & FCCA	NA	NO_COUNTRY_SET
8	W	NULL1 & WORLD	AL	ALBANIA
12	W	NULL1 & WORLD	DZ	ALGERIA

Country-code	Region-code	Domain	ISO-name	Name
16	A	FCC3 & FCCA	AS	AMERICAN SAMOA
...

Creating a FortiAP profile

A FortiAP profile defines radio settings for a particular platform (FortiAP model). The profile also selects which SSIDs (virtual APs) the APs will carry. FortiAP units contain two radio transceivers, making it possible, for example, to provide both 2.4GHz 802.11b/g/n and 5GHz 802.11a/n service from the same access point. The radios can also be used for monitoring, used for the Rogue AP detection feature.

You can modify existing FortiAP profiles or create new ones of your own.



On FortiGate model 30D, web-based manager configuration of FortiAP Profiles is disabled by default. To enable AP profiles, enter the following CLI commands:

```
config system settings
    set gui-ap-profile enable
end
```

To configure a FortiAP profile - web-based manager

1. Go to **WiFi & Switch Controller > FortiAP Profiles** and select **Create New**.
2. Enter a **Name** for the FortiAP Profile.
3. In **Platform**, select the FortiWiFi or FortiAP model to which this profile applies.
4. If split tunneling is used, in **Split Tunneling Subnets**, enter a comma-separated list all of the destination IP address ranges that should **not** be routed through the the FortiGate WiFi controller.
5. For each radio, enter:

Mode	Select the type of mode. Disable – radio disabled Access Point – the platform is an access point Dedicated Monitor – the platform is a dedicated monitor. See Wireless network monitoring on page 1145 .
WIDS Profile	Optionally, select a Wireless Intrusion Detection (WIDS) profile. See Protecting the WiFi network on page 1140 .
Radio Resource Provision	Select to enable the radio resource provision feature. This feature measures utilization and interference on the available channels and selects the clearest channel at each access point. The measurement can be repeated periodically to respond to changing conditions.
Client Load Balancing	Select Frequency Handoff or AP Handoff as needed. See Access point deployment on page 1089 .

Band	<p>Select the wireless protocols that you want to support. The available choices depend on the radio's capabilities. Where multiple protocols are supported, the letter suffixes are combined: "802.11g/b" means 802.11g and 802.11b.</p> <p>Note that on two-radio units such as the FortiAP-221C it is not possible to put both radios on the same band.</p>
Channel Width	Select channel width for 802.11ac or 802.11n on 5GHz.
Short Guard Interval	Select to enable the short guard interval for 802.11ac or 802.11n on 5GHz.
Channels	Select the channel or channels to include. The available channels depend on which IEEE wireless protocol you selected in Band . By default, all available channels are enabled.
TX Power Control	Enable automatic or manual adjustment of transmit power, specifying either minimum and maximum power levels in dBm or as a percentage.
TX Power	<p>When TX Power Control is set to Auto, the TX Power is set by default to a range of 10-17 dBm. Set the range between 1-20 for both the lower and upper limits.</p> <p>When TX Power Control is set to Manual, the TX Power is set by default to 100% of the maximum power permitted in your region. To change the level, drag the slider.</p>
SSIDs	<p>Select between Auto or Manual. Selecting Auto eliminates the need to re-edit the profile when new SSIDs are created. However, you can still select SSIDs individually using Manual.</p> <p>Note that automatic assignment of SSIDs (Auto) is not available for FortiAPs in Local Bridge mode. The option is hidden on both the Managed FortiAP settings and the FortiAP Profile assigned to that AP.</p>

Radio 1 settings are the same as Radio 2 settings except for the options for **Channel**.

Radio 2 settings are available only for FortiAP models with dual radios.

6. Select **OK**.

To configure a FortiAP profile - CLI

This example configures a FortiAP-220B to carry all SSIDs on Radio 1 but only SSID example_wlan on Radio 2.

```
config wireless-controller wtp-profile
  edit guest_prof
    config platform
      set type 220B
    end
    config radio-1
      set mode ap
      set band 802.11g
      set vap-all enable
    end
```

```

config radio-2
    set mode ap
    set band 802.11g
    set vaps example_wlan
end
end

```

Defining a wireless network interface (SSID)

You begin configuring your wireless network by defining one or more SSIDs to which your users will connect. When you create an SSID, a virtual network interface is also created with the **Name** you specified in the SSID configuration. You can configure the settings of an existing SSID in either **WiFi & Switch Controller > SSID** or **System > Network > Interfaces**.



If a software switch interface contains an SSID (but only one), the WiFi SSID settings are available in the switch interface settings.

To create a new SSID

1. Go to **WiFi & Switch Controller > SSID** and select **Create New > SSID**.
2. Fill in the SSID fields as described below.

To configure the settings of an existing SSID

1. Either
 - Go to **WiFi & Switch Controller > SSID**.
 or
 - Go to **Network > Interfaces**.
WiFi interfaces list the SSID beside the interface **Name**.
2. Edit a WiFi interface, modifying the SSID fields as needed.

SSID fields

Interface Name	Enter a name for the SSID interface.
Type	WiFi SSID.
Traffic Mode	<p>Tunnel to Wireless Controller — Data for WLAN passes through WiFi Controller. This is the default.</p> <p>Local bridge with FortiAP's Interface — FortiAP unit Ethernet and WiFi interfaces are bridged.</p> <p>Mesh Downlink — Radio receives data for WLAN from mesh backhaul SSID.</p>
IP/Network Mask	Enter the IP address and netmask for the SSID.

IPv6 Address	Enter the IPv6 address. This is available only when IPv6 has been enabled on the unit.
Administrative Access	Select which types of administrative access are permitted on this SSID.
IPv6 Administrative Access	If you have IPv6 addresses, select the permitted IPv6 administrative access types for this SSID.
DHCP Server	<p>To assign IP addresses to clients, enable DHCP server. You can define IP address ranges for a DHCP server on the FortiGate unit or relay DHCP requests to an external server.</p> <p>If the unit is in transparent mode, the DHCP server settings will be unavailable.</p> <p>For more information, see Configuring DHCP for WiFi clients on page 1073.</p>
Device Detection	Detect connected device type. Enabled by default.
Active Scanning	Enabled by default.
WiFi Settings	
SSID	Enter the SSID. By default, this field contains <code>fortinet</code> .
Security Mode	<p>Select the security mode for the wireless interface. Wireless users must use the same security mode to be able to connect to this wireless interface. Additional security mode options are available in the CLI. For more information, see Configuring security on page 1074.</p> <p>Captive Portal – authenticates users through a customizable web page.</p> <p>WPA2-Personal – WPA2 is WiFi Protected Access version 2. There is one pre-shared key (password) that all users use.</p> <p>WPA2-Personal with Captive Portal – The user will need to know the pre-shared key and will also be authenticated through the custom portal.</p> <p>WPA2-Enterprise – similar to WPA2-Personal, but is best used for enterprise networks. Each user is separately authenticated by user name and password.</p>
Pre-shared Key	Available only when Security Mode is WPA2-Personal . Enter the encryption key that the clients must use.

Authentication	Available only when Security Mode is WPA2-Enterprise . Select one of the following: RADIUS Server — Select the RADIUS server that will authenticate the clients. Local – Select the user group(s) that can authenticate.
Portal Type	Available only when Security Mode is Captive Portal . Choose the captive portal type. Authentication is available with or without a usage policy disclaimer notice.
Authentication Portal	Local - portal hosted on the FortiGate unit External - enter FQDN or IP address of external portal
User Groups	Select permitted user groups for captive portal authentication.
Exempt List	Select exempt lists whose members will not be subject to captive portal authentication.
Customize Portal Messages	Click the listed portal pages to edit them.
Redirect after Captive Portal	Optionally, select Specific URL and enter a URL for user redirection after captive portal authentication. By default, users are redirected to the URL that they originally requested.
Allow New WiFi Client Connections When Controller Is Down	This option is available for local bridge SSIDs with WPA-Personal security. See Combining WiFi and wired networks with a software switch on page 1122 .
Broadcast SSID	Optionally, disable broadcast of SSID. By default, the SSID is broadcast. For more information, see Introduction to wireless networking on page 1048 .
Schedule	Select when the SSID is enabled. You can choose any schedule defined in Policy & Objects > Objects > Schedules .
Block Intra-SSID Traffic	Select to enable the unit to block intra-SSID traffic.
Maximum Clients	Select to limit the number of clients permitted to connect simultaneously. Enter the limit value.
Split Tunneling	Select to enable some subnets to remain local to the remote FortiAP. Traffic for these networks is not routed through the WiFi Controller. Specify split-tunnel networks in the FortiAP Profile. See Split tunneling on page 1128 .
Optional VLAN ID	Enter the ID of the VLAN this SSID belongs to. Enter 0 for non-VLAN operation.

Enable Explicit Web Proxy	Select to enable explicit web proxy for the SSID.
Listen for RADIUS Accounting Messages	Enable if you are using RADIUS-based single sign-on (SSO).
Secondary IP Address	Optionally, enable and define secondary IP addresses. Administrative access can be enabled on secondary interfaces.
Comments	Enter a description or comment for the SSID.

To configure a virtual access point (SSID) - CLI

The example below creates an access point with SSID “example” and WPA2-Personal security. The wireless interface is named example_wlan.

WiFi SSIDs include a schedule that determines when the WiFi network is available. The default schedule is Always. You can choose any schedule (but not schedule group) that is defined in **Policy & Objects > Objects > Schedules**.

```
config wireless-controller vap
  edit example_wlan
    set ssid "example"
    set broadcast-ssid enable
    set security wpa2-only-personal
    set passphrase "hardtoguess"
    set schedule always
    set vdom root
  end
config system interface
  edit example_wlan
    set ip 10.10.120.1 255.255.255.0
  end
```

Configuring DHCP for WiFi clients

Wireless clients need to have IP addresses. If you use RADIUS authentication, each user’s IP address can be stored in the Framed-IP-Address attribute. Otherwise, you need to configure a DHCP server on the WLAN interface to assign IP addresses to wireless clients.

To configure a DHCP server for WiFi clients - web-based manager

1. Go to **WiFi & Switch Controller > SSID** and edit your SSID entry.
2. In **DHCP Server** select **Enable**.
3. In **Address Range**, select **Create New**.
4. In the **Starting IP** and **End IP** fields, enter the IP address range to assign.
By default an address range is created in the same subnet as the wireless interface IP address, but not including that address.
5. Set the **Netmask** to an appropriate value, such as 255.255.255.0.
6. Set the **Default Gateway** to **Same as Interface IP**.

7. Set the **DNS Server** to **Same as System DNS**.
8. If you want to restrict access to the wireless network by MAC address, see [Adding a MAC filter on page 1076](#).
9. Select **OK**.

To configure a DHCP server for WiFi clients - CLI

In this example, WiFi clients on the example_wlan interface are assigned addresses in the 10.10.120.2-9 range to connect with the WiFi access point on 10.10.120.1.

```
config system dhcp server
  edit 0
    set default-gateway 10.10.120.1
    set dns-service default
    set interface example_wlan
    set netmask 255.255.255.0
    config ip-range
      edit 1
        set end-ip 10.10.120.9
        set start-ip 10.10.120.2
      end
    end
  end
```



You cannot delete an SSID (wireless interface) that has DHCP enabled on it.

Configuring security

Using the web-based manager, you can configure captive portal security or WiFi Protected Access version 2 (WPA2) security modes WPA2-Personal and WPA2-Enterprise. Using the CLI, you can also choose WPA/WPA2 modes that support both WPA version 1 and WPA version 2.

WPA2 security with a pre-shared key for authentication is called WPA2-Personal. This can work well for one person or a small group of trusted people. But, as the number of users increases, it is difficult to distribute new keys securely and there is increased risk that the key could fall into the wrong hands.

A more secure form of WPA2 security is WPA2-Enterprise. Users each have their own authentication credentials, verified through an authentication server, usually RADIUS. FortiOS can also authenticate WPA2-Enterprise users through its built-in user group functionality. FortiGate user groups can include RADIUS servers and can select users by RADIUS user group. This makes possible Role-Based Access Control (RBAC).

By default, WPA2 security encrypts communication using Advanced Encryption Standard (AES). But some older wireless clients support only Temporal Key Integrity Protocol (TKIP). You can change the encryption to TKIP or negotiable TKIP-AES in the CLI. For example, to accommodate clients with either TKIP or AES, enter:

```
config wireless-controller vap
  edit example_wlan
    set security wpa-personal
    set passphrase "hardtoguess"
    set encrypt TKIP-AES
  end
```

Captive portal security connects users to an open web portal defined in replacement messages. To navigate to any location beyond the web portal, the user must pass FortiGate user authentication.

WPA-Personal security

WPA2-Personal security setup requires only the preshared key that you will provide to your clients.

To configure WPA2-Personal security - web-based manager

1. Go to **WiFi & Switch Controller > SSID** and edit your SSID entry.
2. In **Security Mode**, select **WPA2 Personal**.
3. In **Pre-shared Key**, enter a key between 8 and 63 characters long.
4. Select **OK**.

To configure WPA2-Personal security - CLI

```
config wireless-controller vap
  edit example_wlan
    set security wpa2-personal
    set passphrase "hardtoguess"
  end
```

WPA-Enterprise security

If you will use FortiOS user groups for authentication, go to **User & Device > User > User Groups** and create those groups first. The groups should be Firewall groups.

If you will use a RADIUS server to authenticate wireless clients, you must first configure the FortiGate unit to access the RADIUS server.

To configure FortiGate unit access to the RADIUS server - web-based manager

1. Go to **User & Device > RADIUS Servers** and select **Create New**.
2. Enter a **Name** for the server.
3. In **Primary Server Name/IP**, enter the network name or IP address for the server.
4. In **Primary Server Secret**, enter the shared secret used to access the server.
5. Optionally, enter the information for a secondary or backup RADIUS server.
6. Select **OK**.

To configure the FortiGate unit to access the RADIUS server - CLI

```
config user radius
  edit exampleRADIUS
    set auth-type auto
    set server 10.11.102.100
    set secret aoewmntiasf
  end
```

RADIUS Change of Authorization (CoA) support

The CoA feature enables the FortiGate to receive a client disconnect message from the RADIUS server. This is used to disconnect clients when their time, credit or bandwidth had been used up. Enable this on the RADIUS server using the CLI:

```
config user radius
  edit <name>
```

```
set radius-coa enable
end
```

To configure WPA-Enterprise security - web-based manager

1. Go to **WiFi & Switch Controller > SSID** and edit your SSID entry.
2. In **Security Mode**, select **WPA2 Enterprise**.
3. In **Authentication**, do one of the following:
 - If you will use a RADIUS server for authentication, select **RADIUS** Server and then select the RADIUS server.
 - If you will use a local user group for authentication, select **Local** and then select the user group(s) permitted to use the wireless network.
4. Select **OK**.

To configure WPA-Enterprise security - CLI

```
config wireless-controller vap
edit example_wlan
set security wpa2-enterprise
set auth radius
set radius-server exampleRADIUS
end
```

Captive portal security

Captive portal security provides an access point that initially appears open. The wireless client can connect to the AP with no security credentials. The AP responds to the client's first HTTP request with a web page requesting user name and password. Until the user enters valid credentials, no communication beyond the AP is permitted.

The captive portal can be hosted on the FortiGate unit, or externally. For details see

[Configuring WiFi captive portal security - FortiGate captive portal on page 1078](#)

[Configuring WiFi captive portal security - external server on page 1078](#)

For general information about captive portals, see the Captive Portal chapter of the Authentication Guide.

Adding a MAC filter

On each SSID, you can create a MAC address filter list to either permit or exclude a list of clients identified by their MAC addresses.

This is actually not as secure as it appears. Someone seeking unauthorized access to your network can obtain MAC addresses from wireless traffic and use them to impersonate legitimate users. A MAC filter list should only be used in conjunction with other security measures such as encryption.

To configure a MAC filter - web-based manager

1. Go to **WiFi & Switch Controller > SSID** and edit your SSID entry.
2. In the **DHCP Server** section, expand **Advanced**.
3. In **MAC Reservation + Access Control**, double-click in the **Unknown MAC Addresses** line and select **Assign IP** or **Block**, as needed.
By default, unlisted MAC addresses are assigned an IP address automatically.
4. In **MAC Reservation + Access Control**, select **Create New**.

5. Enter a MAC address in the **MAC** field.
6. In **IP or Action**, select one of:
 - **Reserve IP** — enter the IP address that is always assigned to this MAC address.
 - **Assign IP** — an IP address is assigned to this MAC address automatically.
 - **Block** — This MAC address will not be assigned an IP address.
7. Repeat steps 4 through 6 for each additional MAC address that you want to add.
8. Select **OK**.

To configure a MAC filter - CLI

1. Enter

```
config system dhcp server
show
```
2. Find the entry where `interface` is your WiFi interface. Edit that entry and configure the MAC filter. In this example, the MAC address 11:11:11:11:11:11 will be excluded. Unlisted MAC addresses will be assigned an IP address automatically.

```
edit 3
config reserved-address
edit 1
set action block
set mac 11:11:11:11:11:11
end
set mac-acl-default-action assign
end
```

Limiting the number of clients

You might want to prevent overloading of your access point by limiting the number of clients who can associate with it at the same time. Limits can be applied per SSID, per AP, or per radio.

To limit the number of clients per SSID - GUI

1. Go to **WiFi & Switch Controller > SSID** and edit your SSID.
2. Turn on **Maximum Clients** and enter the maximum number of clients in **Limit Concurrent WiFi Clients**.

To limit the number of clients per AP - CLI

Edit the wtp-profile (FortiAP profile), like this:

```
config wireless-controller wtp-profile
edit "FAP221C-default"
set max-clients 30
end
```

To limit the number of clients per radio - CLI

Edit the wtp-profile (FortiAP profile), like this:

```
config wireless-controller wtp-profile
edit "FAP221C-default"
config radio-1
set max-clients 10
```

```

end
config radio-2
    set max-clients 30
end
end

```

Multicast enhancement

FortiOS can translate multicast traffic into unicast traffic to send to clients, maintaining its own multicast client through IGMP snooping. You can configure this in the CLI:

```

config wireless-controller vap
    edit example_wlan
        set multicast-enhance enable
        set me-disable-thresh 32
    end
end

```

If the number of clients on the SSID is larger than `me-disable-thresh`, multicast enhancement is disabled.

Configuring WiFi captive portal security - FortiGate captive portal

The built-in FortiGate captive portal is simpler than an external portal. It can even be customized if needed.

To configure a WiFi Captive Portal - web-based manager:

1. Go to **WiFi & Switch Controller > SSID** and create your SSID.
If the SSID already exists, you can edit the SSID or you can edit the WiFi interface in **Network > Interfaces**.
2. In **Security Mode**, select **Captive Portal**.
3. Enter

Portal Type	The portal can provide authentication and/or disclaimer, or perform user email address collection. See Defining a wireless network interface (SSID) on page 1070 .
Authentication Portal	Local
User Groups	Select permitted user groups or select Use Groups from Policies , which permits the groups specified in the security policy.
Exempt List	Select exempt lists whose members will not be subject to captive portal authentication.
Customize Portal Messages	Click the link of the portal page that you want to modify. For more information see the Captive Portal chapter of the Authentication Guide.

4. Select **OK**.

Configuring WiFi captive portal security - external server

An external captive portal is a web page on a web server. The essential part of the web portal page is a script that gathers the user's logon credentials and sends back to the FortiGate a specifically-formatted POST message.

The portal page can also contain links to local information such as legal notices, terms of service and so on.

Without authenticating, the user cannot access any other information. This is sometimes called a “walled garden”.

On the captive portal page, the user submits credentials, which the script returns to the FortiGate at the URL **https://<FGT_IP>:1000/fgtauth** with data

magic=session_id&username=<username>&password=<password>.

(The magic value was provided in the initial FortiGate request to the web server.)

To ensure that credentials are communicated securely, enable the use of HTTPS for authentication:

```
config user setting
    set auth-secure-http enable
end
```

To configure use of an external WiFi Captive Portal - web-based manager:

1. Go to **WiFi & Switch Controller > SSID** and create your SSID.
If the SSID already exists, you can edit the SSID or you can edit the WiFi interface in **Network > Interfaces**.
2. In **Security Mode**, select **Captive Portal**.
3. Enter

Portal Type	The portal can provide authentication and/or disclaimer, or perform user email address collection.
Authentication Portal	External - enter the FQDN or IP address of the external portal. Typically, this is the URL of a script. Do not include the protocol (http:// or https://) part of the URL.
User Groups	Select permitted user groups or select Use Groups from Policies , which permits the groups specified in the security policy.
Exempt List	Select exempt lists whose members will not be subject to captive portal authentication.
Redirect after Captive Portal	Original Request Specific URL - enter URL

4. Select **OK**.

Defining SSID groups

Optionally, you can define SSID groups. An SSID group has SSIDs as members and can be specified just like an SSID in a FortiAP Profile.

To create an SSID group - GUI

Go to **WiFi & Switch Controller > SSID** and select **Create New > SSID Group**. Give the group a **Name** and choose **Members** (SSIDs, but not SSID groups).

To create an SSID group - CLI:

```
config wireless-controller vap-group
```

```
edit vap-group-name
  set vaps "ssid1" "ssid2"
end
```

Dynamic user VLAN assignment

Clients connecting to the WiFi network can be assigned to a VLAN. You can do this with RADIUS attributes when the user authenticates or with VLAN pooling when the client associates with a particular FortiAP. You cannot use both of these methods at the same time.

VLAN assignment by RADIUS

You can assign each individual user to a VLAN based on information stored in the RADIUS authentication server. If the user's RADIUS record does not specify a VLAN ID, the user is assigned to the default VLAN for the SSID.

The RADIUS user attributes used for the VLAN ID assignment are:

IETF 64 (Tunnel Type)—Set this to VLAN.

IETF 65 (Tunnel Medium Type)—Set this to 802

IETF 81 (Tunnel Private Group ID)—Set this to the VLAN ID.

To configure dynamic VLAN assignment, you need to:

1. Configure access to the RADIUS server.
2. Create the SSID and enable dynamic VLAN assignment.
3. Create a FortiAP Profile and add the local bridge mode SSID to it.
4. Create the VLAN interfaces and their DHCP servers.
5. Create security policies to allow communication from the VLAN interfaces to the Internet.
6. Authorize the FortiAP unit and assign the FortiAP Profile to it.

To configure access to the RADIUS server

1. Go to **User & Device > RADIUS Servers** and select **Create New**.
2. Enter a **Name**, the name or IP address in **Primary Server IP/Name**, and the server secret in **Primary Server Secret**.
3. Select **OK**.

To create the dynamic VLAN SSID

1. Go to **WiFi & Switch Controller > SSID**, select **Create New > SSID** and enter:

Name	An identifier, such as dynamic_vlan_ssid.
Traffic Mode	Local bridge or Tunnel, as needed.
SSID	An identifier, such as DYNSSID.
Security Mode	WPA2 Enterprise
Authentication	RADIUS Server. Select the RADIUS server that you configured.

2. Select **OK**.
3. Enable dynamic VLAN in the CLI. Optionally, you can also assign a VLAN ID to set the default VLAN for users without a VLAN assignment.

```
config wireless-controller vap
  edit dynamic_vlan_ssid
    set dynamic-vlan enable
    set vlanid 10
  end
```

To create the FortiAP profile for the dynamic VLAN SSID

1. Go to **WiFi & Switch Controller > FortiAP Profiles**, select **Create New** and enter:

Name	A name for the profile, such as dyn_vlan_profile.
Platform	The FortiAP model you are using. If you use more than one model of FortiAP, you will need a FortiAP Profile for each model.
Radio 1 and Radio 2	
SSID	Select the SSID you created (example dynamic_vlan_ssid). Do not add other SSIDs.

2. Adjust other radio settings as needed.
3. Select **OK**.

To create the VLAN interfaces

1. Go to **Network > Interfaces** and select **Create New > Interface**.
2. Enter:

Name	A name for the VLAN interface, such as VLAN100.
Interface	The physical interface associated with the VLAN interface.
VLAN ID	The numeric VLAN ID, for example 100.
Addressing mode	Select Manual and enter the IP address / Network Mask for the virtual interface.
DHCP Server	Enable and then select Create New to create an address range.

3. Select **OK**.
4. Repeat the preceding steps to create other VLANs as needed.

Security policies determine which VLANs can communicate with which other interfaces. These are the simple Firewall Address policy without authentication. Users are assigned to the appropriate VLAN when they authenticate.

To connect and authorize the FortiAP unit

1. Connect the FortiAP unit to the FortiGate unit.
2. Go to **WiFi & Switch Controller > Managed FortiAPs**.

3. When the FortiAP unit is listed, double-click the entry to edit it.
4. In **FortiAP Profile**, select the FortiAP Profile that you created.
5. Select **Authorize**.
6. Select **OK**.

VLAN assignment by VLAN pool

In an SSID, you can define a VLAN pool. As clients associate to an AP, they are assigned to a VLAN. A VLAN pool can

- assign a specific VLAN based on the AP's FortiAP group, usually for network configuration reasons, or
- assign one of several available VLANs for network load balancing purposes (tunnel mode SSIDs only)

To assign a VLAN by FortiAP group - CLI

In this example, VLAN 101, 102, or 103 is assigned depending on the AP's FortiAP group.

```
config wireless-controller vap
  edit wlan
    set vlan-pooling wtp-group
    config vlan-pool
      edit 101
        set wtp-group wtpgrp1
      next
      edit 102
        set wtp-group wtpgrp2
      next
      edit 103
        set wtp-group wtpgrp3
      end
    end
  end
end
```

Load balancing

There are two VLAN pooling methods used for load balancing:

The choice of VLAN can be based on any one of the following criteria:

- **round-robin** - from the VLAN pool, choose the VLAN with the smallest number of clients
- **hash** - choose a VLAN from the VLAN pool based on a hash of the current number of SSID clients and the number of entries in the VLAN pool

If the VLAN pool contains no valid VLAN ID, the SSID's static VLAN ID setting is used.

To assign a VLAN by round-robin selection - CLI

In this example, VLAN 101, 102, or 103 is assigned using the round-robin method:

```
config wireless-controller vap
  edit wlan
    set vlan-pooling round-robin
    config vlan-pool
      edit 101
      next
      edit 102
```

```
        next
        edit 103
        end
    end
end
```

To assign a VLAN by hash-based selection - CLI

In this example, VLAN 101, 102, or 103 is assigned using the hash method:

```
config wireless-controller vap
edit wlan
set vlan-pooling hash
config vlan-pool
edit 101
next
edit 102
next
edit 103
end
end
end
```

Configuring user authentication

You can perform user authentication when the wireless client joins the wireless network and when the wireless user communicates with another network through a firewall policy. WEP and WPA-Personal security rely on legitimate users knowing the correct key or passphrase for the wireless network. The more users you have, the more likely it is that the key or passphrase will become known to unauthorized people. WPA-Enterprise and captive portal security provide separate credentials for each user. User accounts can be managed through FortiGate user groups or an external RADIUS authentication server.

WPA2 Enterprise authentication

Enterprise authentication can be based on the local FortiGate user database or on a remote RADIUS server. Local authentication is essentially the same for WiFi users as it is for wired users, except that authentication for WiFi users occurs when they associate their device with the AP. Therefore, enterprise authentication must be configured in the SSID. WiFi users can belong to user groups just the same as wired users and security policies will determine which network services they can access.

If your WiFi network uses WPA2 Enterprise authentication verified by a RADIUS server, you need to configure the FortiGate unit to connect to that RADIUS server.

Configuring connection to a RADIUS server - web-based manager

1. Go to **User & Device > RADIUS Servers** and select **Create New**.
2. Enter a **Name** for the server.
This name is used in FortiGate configurations. It is not the actual name of the server.
3. In **Primary Server Name/IP**, enter the network name or IP address for the server.
4. In **Primary Server Secret**, enter the shared secret used to access the server.
5. Optionally, enter the information for a secondary or backup RADIUS server.
6. Select **OK**.

To configure the FortiGate unit to access the RADIUS server - CLI

```
config user radius
  edit exampleRADIUS
    set auth-type auto
    set server 10.11.102.100
    set secret aoewmntiasf
  end
```

To implement WPA2 Enterprise security, you select this server in the SSID security settings. See [Configuring user authentication on page 1083](#).

To use the RADIUS server for authentication, you can create individual FortiGate user accounts that specify the authentication server instead of a password, and you then add those accounts to a user group. Or, you can add the authentication server to a FortiGate user group, making all accounts on that server members of the user group.

Creating a wireless user group

Most wireless networks require authenticated access. To enable creation of firewall policies specific to WiFi users, you should create at least one WiFi user group. You can add or remove users later. There are two types of user group to consider:

- A Firewall user group can contain user accounts stored on the FortiGate unit or external authentication servers such as RADIUS that contain and verify user credentials.
- A Fortinet single sign-on (FSSO) user group is used for integration with Windows Active Directory or Novell eDirectory. The group can contain Windows or Novell user groups who will be permitted access to the wireless LAN.

WiFi single sign-on (WSSO) authentication

WSSO is RADIUS-based authentication that passes the user's user group memberships to the FortiGate. For each user, the RADIUS server must provide user group information in the Fortinet-Group-Name attribute. This information is stored in the server's database. After the user authenticates, security policies provide access to network services based on user groups.

1. Configure the RADIUS server to return the Fortinet-Group-Name attribute for each user.
2. Configure the FortiGate to access the RADIUS server, as described in [WPA2 Enterprise authentication on page 1083](#).
3. Create firewall user groups on the FortiGate with the same names as the user groups listed in the RADIUS database. Leave the groups empty.
4. In the SSID choose WPA2-Enterprise authentication. In the **Authentication** field, select **RADIUS Server** and choose the RADIUS server that you configured.
5. Create security policies as needed, using user groups (**Source User(s)** field) to control access.

When a user authenticates by WSSO, the firewall monitor **Monitor > Firewall User Monitor**) shows the authentication method as WSSO.

Assigning WiFi users to VLANs dynamically

Some enterprise networks use Virtual LANs (VLANs) to separate traffic. In this environment, to extend network access to WiFi users might appear to require multiple SSIDs. But it is possible to automatically assign each user to their appropriate VLAN from a single SSID. To accomplish this requires RADIUS authentication that passes the appropriate VLAN ID to the FortiGate by RADIUS attributes. Each user's VLAN assignment is stored in the user database of the RADIUS server.

1. Configure the RADIUS server to return the following attributes for each user:
Tunnel-Type (value: VLAN)
Tunnel-Medium-Type (value: IEEE-802)
Tunnel_Private-Group-Id (value: the VLAN ID for the user's VLAN)
2. Configure the FortiGate to access the RADIUS server.
3. Configure the SSID with WPA2-Enterprise authentication. In the **Authentication** field, select **RADIUS Server** and choose the RADIUS server that you will use.
4. Create VLAN subinterfaces on the SSID interface, one for each VLAN. Set the VLAN ID of each as appropriate. You can do this on the **Network > Interfaces** page.
5. Enable Dynamic VLAN assignment for the SSID. For example, if the SSID interface is "office", enter:

```
config wireless-controller vap
  edit office
    set dynamic-vlan enable
  end
```
6. Create security policies for each VLAN. These policies have a WiFi VLAN subinterface as **Incoming Interface** and allow traffic to flow to whichever **Outgoing Interface** these VLAN users will be allowed to access.

MAC-based authentication

Wireless clients can also be supplementally authenticated by MAC address. A RADIUS server stores the allowed MAC address for each client and the wireless controller checks the MAC address independently of other authentication methods.

MAC-based authentication must be configured in the CLI. In the following example, MAC-based authentication is added to an existing access point "vap1" to use RADIUS server hq_radius (configured on the FortiGate):

```
config wireless-controller vap
  edit vap1
    set radius-mac-auth enable
    set radius-mac-auth-server hq_radius
  end
```

Authenticating guest WiFi users

The FortiOS Guest Management feature enables you to easily add guest accounts to your FortiGate unit. These accounts are authenticate guest WiFi users for temporary access to a WiFi network managed by a FortiGate unit. To implement guest access, you need to

1. Go to **User & Device > User Groups** and create one or more guest user groups.
2. Go to **User & Device > Guest Management** to create guest accounts. You can print the guest account credentials or send them to the user as an email or SMS message.
3. Go to **WiFi & Switch Controller > SSID** and configure your WiFi SSID to use captive portal authentication. Select the guest user group(s) that you created.

Guest users can log into the WiFi captive portal with their guest account credentials until the account expires. For more detailed information about creating guest accounts, see "Managing Guest Access" in the Authentication chapter of the FortiOS Handbook.

Configuring firewall policies for the SSID

For users on the WiFi LAN to communicate with other networks, firewall policies are required. This section describes creating a WiFi network to Internet policy.

Before you create firewall policies, you need to define any firewall addresses you will need.

To create a firewall address for WiFi users - web-based manager

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**, enter the following information and select **OK**.

Name	Enter a name for the address, wifi_net for example.
Type	Select Subnet .
Subnet / IP Range	Enter the subnet address, 10.10.110.0/24 for example.
Interface	Select the interface where this address is used, e.g., example_wifi

To create a firewall address for WiFi users - CLI

```
config firewall address
  edit "wifi_net"
    set associated-interface "example_wifi"
    set subnet 10.10.110.0 255.255.255.0
  end
```

To create a firewall policy - web-based manager

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. In **Incoming Interface**, select the wireless interface.
3. In **Source Address**, select the address of your WiFi network, wifi_net for example.
4. In **Outgoing Interface**, select the Internet interface, for example, port1.
5. In **Destination Address**, select **All**.
6. In **Service**, select **ALL**, or select the particular services that you want to allow, and then select the right arrow button to move the service to the **Selected Services** list.
7. In **Schedule**, select **always**, unless you want to define a schedule for limited hours.
8. In **Action**, select **ACCEPT**.
9. Select **Enable NAT**.
10. Optionally, set up UTM features for wireless users.
11. Select **OK**.

To create a firewall policy - CLI

```
config firewall policy
  edit 0
    set srcintf "example_wifi"
    set dstintf "port1"
    set srcaddr "wifi_net"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY"
    set nat enable
  end
```

Configuring the built-in access point on a FortiWiFi unit

Both FortiGate and FortiWiFi units have the WiFi controller feature. If you configure a WiFi network on a FortiWiFi unit, you can also use the built-in wireless capabilities in your WiFi network as one of the access points.

If Virtual Domains are enabled, you must select the VDOM to which the built-in access point belongs. You do this in the CLI. For example:

```
config wireless-controller global
    set local-radio-vdom vdom1
end
```

To configure the FortiWiFi unit's built-in WiFi access point

1. Go to **WiFi Controller > Local WiFi Radio**.
2. Make sure that **Enable WiFi Radio** is selected.
3. In **SSID**, if you do not want this AP to carry all SSIDs, select **Select SSIDs** and then select the required SSIDs.
4. Optionally, adjust the **TX Power** slider.
If you have selected your location correctly (see [Configuring the built-in access point on a FortiWiFi unit on page 1087](#)), the 100% setting corresponds to the maximum power allowed in your region.
5. If you do not want the built-in WiFi radio to be used for rogue scanning, select **Do not participate in Rogue AP scanning**.
6. Select **OK**.

If you want to connect external APs, such as FortiAP units, see the next chapter, [Access point deployment](#).

Enforcing UTM policies on a local bridge SSID for managed smart APs

The `config wireless-controller utm-profile` command lets administrators configure UTM profiles in order to enforce UTM policies on a local bridge SSID when Smart AP's are managed by FortiGate.

As a result, these UTM profiles can also be assigned under `config wireless-controller vap`.

Please note that this is only supported in Bridge-mode.

In addition, a new diagnose command has been introduced to determine the status of the `cw_acd` daemon, which handles the communication between FortiGate and APs.

Note that the default `utm-profile` available (named `wifi-default`) has all applicable options within the command set to `wifi-default`.

Use "?" to view all available profiles to assign, for example, "`set ips-sensor ?`".

Syntax:

```
config wireless-controller utm-profile
    edit <name>
        set comment <comment>
        set utm-log {enable | disable}
        set ips-sensor <name>
        set application-list <name>
        set antivirus-profile <name>
        set webfilter-profile <name>
        set firewall-policy <id>
        set scan-botnet-connections {disable | block | monitor}
```

```
    next
end

config wireless-controller vap
    edit <name>
        set utm-profile <name>
    next
end
```

To debug the `cw_acd_helper` daemon, use the following diagnose command:

```
diagnose wireless-controller wlac_hlp
```

Access point deployment

This chapter describes how to configure access points for your wireless network.

Overview

FortiAP units discover WiFi controllers. The administrator of the WiFi controller authorizes the FortiAP units that the controller will manage.

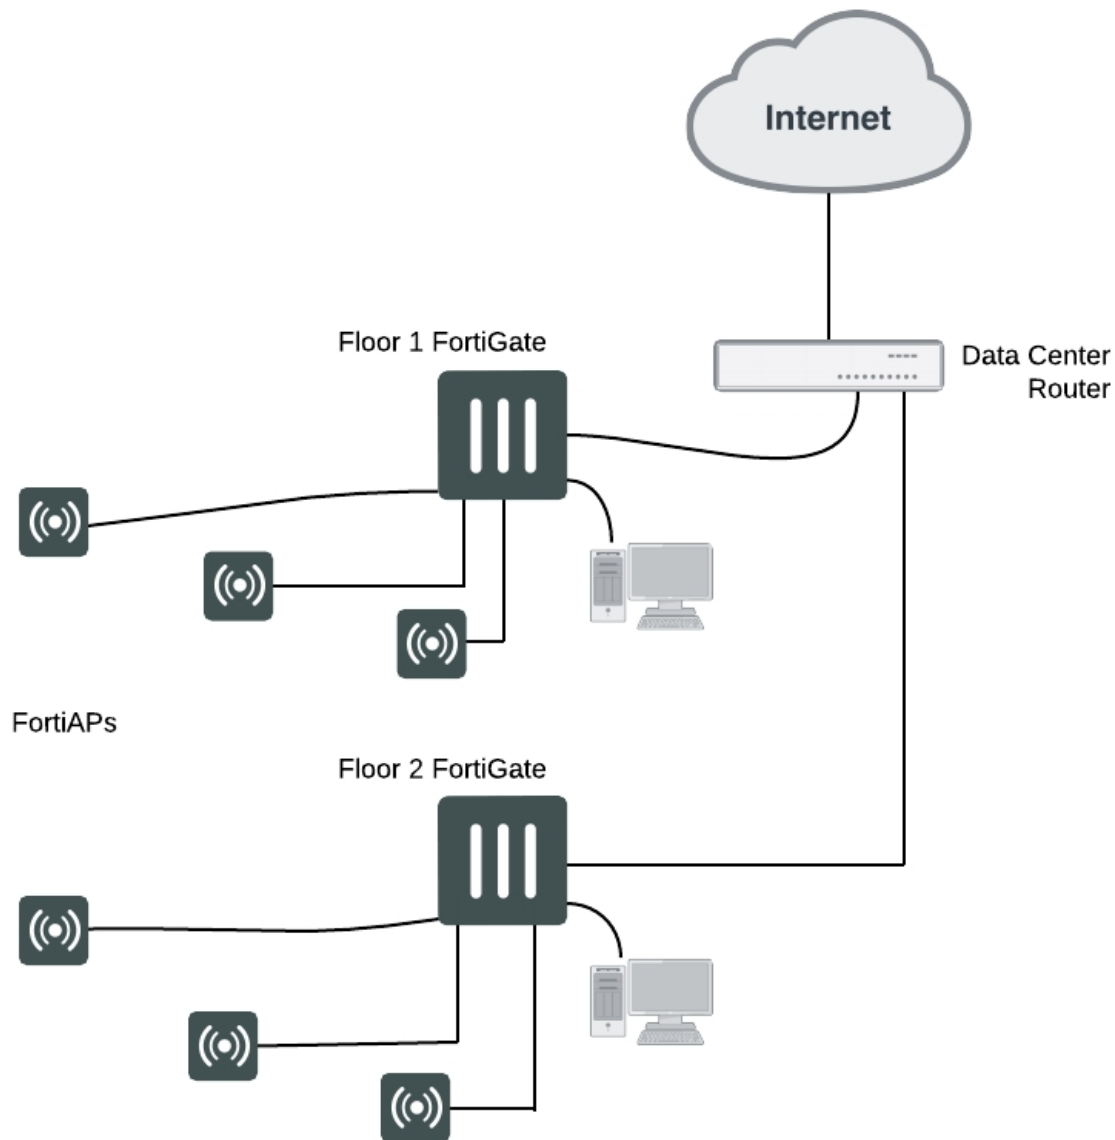
In most cases, FortiAP units can find WiFi controllers through the wired Ethernet without any special configuration. Review the following section, [Access point deployment on page 1089](#), to make sure that your method of connecting the FortiAP unit to the WiFi controller is valid. Then, you are ready to follow the procedures in [Access point deployment on page 1089](#).

If your FortiAP units are unable to find the WiFi controller, refer to [Access point deployment on page 1089](#) for detailed information about the FortiAP unit's controller discovery methods and how you can configure them.

Network topology for managed APs

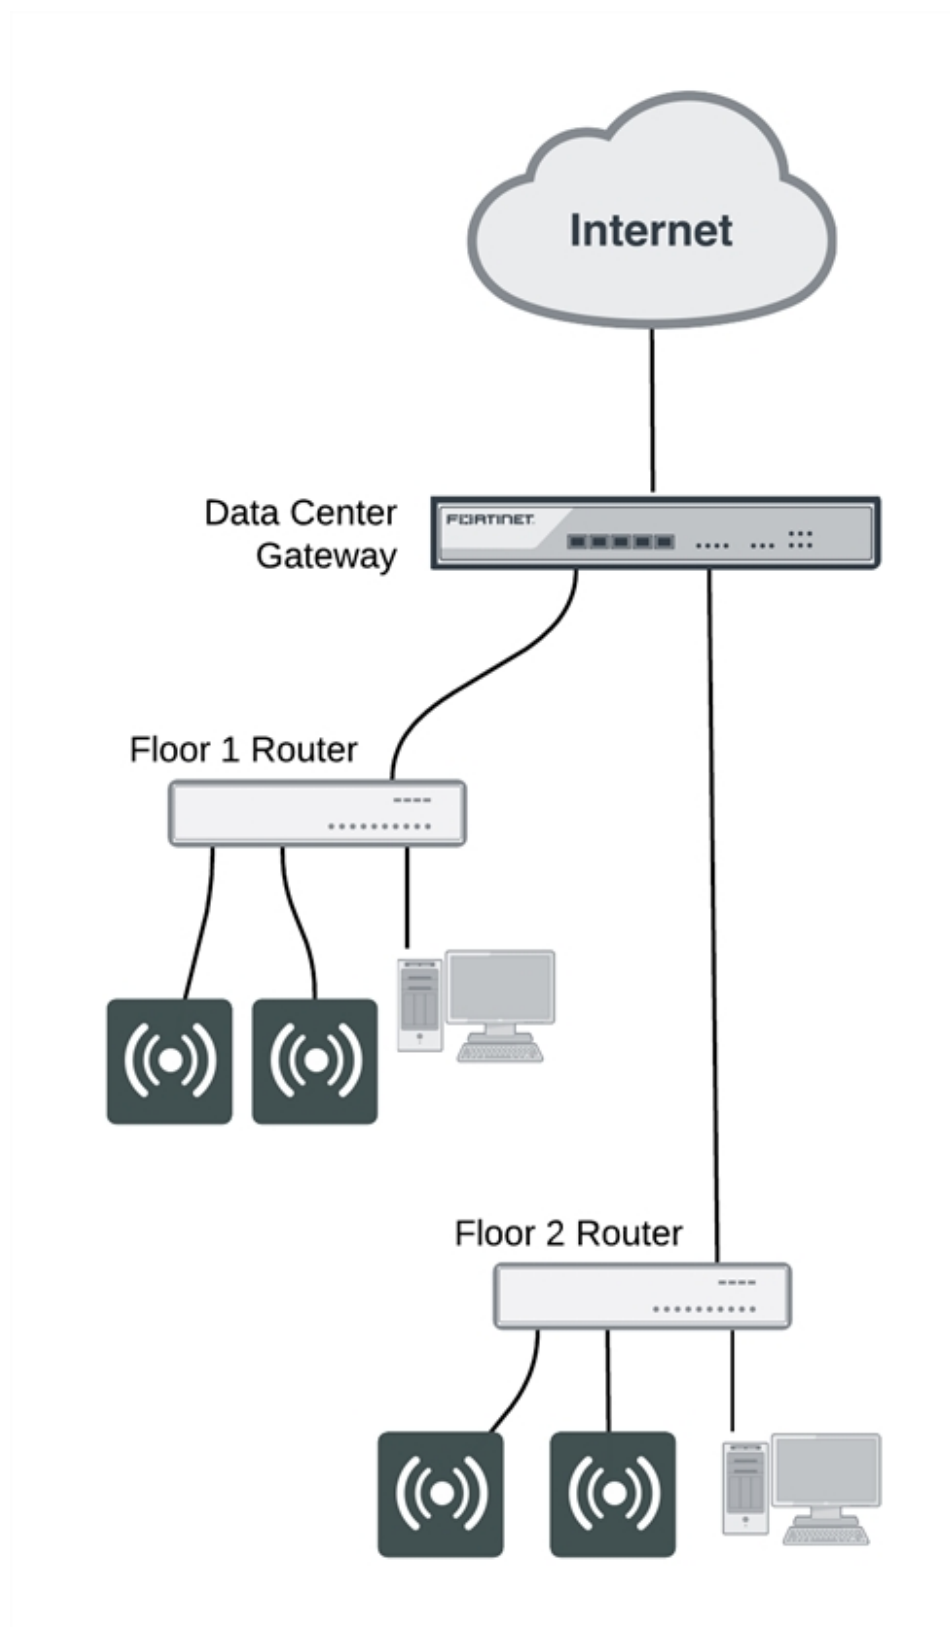
The FortiAP unit can be connected to the FortiGate unit in any of the following ways:

Direct connection: The FortiAP unit is directly connected to the FortiGate unit with no switches between them. This configuration is common for locations where the number of FortiAP's matches up with the number of 'internal' ports available on the FortiGate. In this configuration the FortiAP unit requests an IP address from the FortiGate unit, enters discovery mode and should quickly find the FortiGate WiFi controller. This is also known as a wirecloset deployment. See "Wirecloset and Gateway deployments" below.

Wirecloset deployment


Switched Connection: The FortiAP unit is connected to the FortiGate WiFi controller by an Ethernet switch operating in L2 switching mode or L3 routing mode. There must be a routable path between the FortiAP unit and the FortiGate unit and ports 5246 and 5247 must be open. This is also known as a gateway deployment. See Gateway Deployment below.

Gateway Deployment



Connection over WAN: The FortiGate WiFi controller is off-premises and connected by a VPN tunnel to a local FortiGate. In this method of connectivity its best to configure each FortiAP with the static IP address of the WiFi controller. Each FortiAP can be configured with three WiFi controller IP addresses for redundant failover. This is also known as a datacenter remote management deployment. See Remote deployment below.

Remote deployment



Discovering and authorizing APs

After you prepare your FortiGate, you can connect your APs to discover them using the discovery methods described earlier. To prepare the FortiGate, you need to

- Configure the network interface to which the AP will connect.
- Configure DHCP service on the interface to which the AP will connect.
- Optionally, preauthorize FortiAP units. They will begin to function when connected.
- Connect the AP units and let the FortiGate unit discover them.
- Enable each discovered AP and configure it or assign it to an AP profile.

Configuring the network interface for the AP unit

The interface to which you connect your wireless access point needs an IP address. No administrative access, DNS Query service or authentication should be enabled.

To configure the interface for the AP unit - web-based manager

1. Go to **Network > Interfaces** and edit the interface to which the AP unit connects.
2. Set **Addressing Mode** to **Dedicate to Extension Device**.
3. Enter the IP address and netmask to use.
This FortiGate unit automatically configures a DHCP server on the interface that will assign the remaining higher addresses up to .254 to FortiAP units. For example, if the IP address is 10.10.1.100, the FortiAP units will be assigned 10.10.1.101 to 10.10.1.254. To maximize the available addresses, use the .1 address for the interface: 10.10.1.1, for example.
4. Select **OK**.

To configure the interface for the AP unit - CLI

In the CLI, you must configure the interface IP address and DHCP server separately.

```
config system interface
  edit port3
    set mode static
    set ip 10.10.70.1 255.255.255.0
  end
config system dhcp server
  edit 0
    set interface "dmz"
    config ip-range
      edit 1
        set end-ip 10.10.70.254
        set start-ip 10.10.70.2
      end
    set netmask 255.255.255.0
    set vci-match enable
    set vci-string "FortiAP"
  end
end
```

The optional `vci-match` and `vci-string` fields ensure that the DHCP server will provide IP addresses only to FortiAP units.

Pre-authorizing a FortiAP unit

If you enter the FortiAP unit information in advance, it is authorized and will begin to function when it is connected.

To pre-authorize a FortiAP unit

1. Go to **WiFi & Switch Controller > Managed FortiAPs** and select **Create New**.
On some models the **WiFi Controller** menu is called **WiFi & Switch Controller**.
2. Enter the **Serial Number** of the FortiAP unit.
3. Configure the **Wireless Settings** as required.
4. Select **OK**.

Enabling and configuring a discovered AP

Within two minutes of connecting the AP unit to the FortiGate unit, the discovered unit should be listed on **WiFi Controller > Managed FortiAPs** page. After you select the unit, you can authorize, edit or delete it.

Discovered access point unit

<div> + Create New Edit Delete Refresh Authorize </div> <div> AP Radio Managed FortiAPs 0/32 </div>								
Mesh	Access Point	State	Connected Via	SSIDs	Channel	Clients	OS Version	FortiAP Profile
	FP221C3X14019926	?	192.168.2.2	Radio 1: Radio 2: Student-net	Radio1: 0 Radio2: 0	Radio 1: 0 Radio 2: 0		FAP221C-default

When you authorize (enable) a FortiAP unit, it is configured by default to use the default FortiAP profile (determined by model). You can create and select a different profile if needed. The FortiAP profile defines the entire configuration for the AP.

To add and configure the discovered AP unit - web-based manager

1. Go to **WiFi & Switch Controller > Managed FortiAPs**.
This configuration also applies to local WiFi radio on FortiWiFi models.
2. Select the FortiAP unit from the list and edit it.
3. Optionally, enter a **Name**. Otherwise, the unit will be identified by serial number.
4. Select **Authorize**.
5. Select a **FortiAP Profile**.
6. Select **OK**.

The physical access point is now added to the system. If the rest of the configuration is complete, it should be possible to connect to the wireless network through the AP.

To add the discovered AP unit - CLI

First get a list of the discovered access point unit serial numbers:

```
get wireless-controller wtp
```

Add a discovered unit and associate it with AP-profile1, for example:

```
config wireless-controller wtp
  edit FAP22A3U10600118
    set admin enable
    set wtp-profile AP-profile1
  end
```

To view the status of the added AP unit

```
config wireless-controller wtp
  edit FAP22A3U10600118
  get
```

The `join-time` field should show a time, not "N/A". See the preceding web-based manager procedure for more information.

Disable automatic discovery of unknown FortiAPs

By default, the FortiGate adds newly discovered FortiAPs to the Managed FortiAPs list, awaiting the administrator's authorization. Optionally, you can disable this automatic registration function to avoid adding unknown FortiAPs. A FortiAP will be registered and listed only if its serial number has already been added manually to the Managed FortiAPs list. AP registration is configured on each interface.

To disable automatic discovery and registration, enter the following command:

```
config system interface
  edit port15
```

```
set ap-discover disable
end
```

Automatic authorization of extension devices

To simplify adding FortiAP or FortiSwitch devices to your network, you can enable automatic authorization of devices as they are connected, instead of authorizing each one individually.

This feature is only configurable in the CLI.

To enable automatic authorization on all dedicated interfaces

```
config system global
set auto-auth-extension-device enable
end
```

To enable automatic authorization per-interface

```
config system interface
edit <port>
set auto-auth-extension-device enable
end
```

Assigning the same profile to multiple FortiAP units

The same profile can now be applied to multiple managed FortiAP units at the same time. To do this, do the following:

1. Go to **WiFi & Switch Controller > Managed FortiAPs** to view the AP list.
2. Select all FortiAP units you wish to apply the profile to.
3. Right click on one of the selected FortiAPs and select **Assign Profile**.
4. Choose the profile you wish to apply.

Overriding the FortiAP profile

In the FortiAP configuration **WiFi & Switch Controller > Managed FortiAPs**, there several radio settings under **Override Radio 1** and **Override Radio 2** to choose a value independently of the FortiAP profile setting. When each of the radios are disabled, you will see what the FortiAP Profile has each of the settings configured to.

Band	The available options depend on the capability of the radio. Overriding Band also overrides Channels . Make appropriate settings in Channels .
Channels	Choose channels. The available channels depend on the Band.
TX Power Control	If you enable Auto , adjust to set the power range in dBm. If you enable Manual , adjust the slider. The 100% setting is the maximum power permitted in your region. See Configuring a WiFi LAN on page 1065 .
SSIDs	Select between Auto or Manual . Selecting Auto eliminates the need to re-edit the profile when new SSIDs are created. However, you can still select SSIDs individually using Manual .

To override radio settings in the CLI

In this example, Radio 1 is set to 802.11n on channel 11, regardless of the profile setting.

```
config wireless-controller wtp
edit FP221C3X14019926
config radio-1
set override-band enable
set band 802.11n
set override-channel enable
set channel 11
end
```

Override settings are available for band, channel, vaps (SSIDs), and txpower.

Outside of configuring radio settings, you can also override FortiAP LED state, WAN port mode, IP Fragmentation prevention method, spectrum analysis, split tunneling, and login password settings.

Accessing the FortiAP CLI through the FortiGate unit

Enable remote login for the FortiAP. In the FortiAP Profile for this FortiAP, enable remote access.

Connecting to the FortiAP CLI

The FortiAP unit has a CLI through which some configuration options can be set. You can access the CLI using Telnet.

To access the FortiAP unit CLI through the FortiAP Ethernet port

1. Connect your computer to the FortiAP Ethernet interface, either directly with a cross-over cable or through a separate switch or hub.
2. Change your computer's IP address to 192.168.1.3
3. Telnet to IP address 192.168.1.2.
Ensure that FortiAP is in a private network with no DHCP server for the static IP address to be accessible.
4. Login with user name admin and no password.
5. Enter commands as needed.
6. Optionally, use the `passwd` command to assign an administrative password for better security.
7. Save the configuration by entering the following command:

```
cfg -c .
```
8. Unplug the FortiAP and then plug it back in, in order for the configuration to take effect

Accessing the FortiAP CLI through the FortiGate

After the FortiAP has been installed, physical access to the unit might be inconvenient. You can access a connected FortiAP unit's CLI through the FortiGate unit that controls it.

To enable remote access to the FortiAP CLI

In the CLI, edit the FortiAP Profile that applies to this FortiAP.

```
config wireless-controller wtp-profile
edit FAP221C-default
set allowaccess telnet
end
```




FortiAP now supports HTTPS and SSH administrative access, as well as HTTP and Telnet. Use the command above to set administrative access to `telnet`, `http`, `https`, or `ssh`.

To access the FortiAP unit CLI through the FortiGate unit - GUI

1. Go to **WiFi & Switch Controller > Managed FortiAPs**.
2. In the list, right-click the FortiAP unit and select **>_Connect to CLI**.
A detached Console window opens.
3. At the FortiAP login prompt, enter `admin`. When you are finished using the FortiAP CLI, enter `exit`.

To access the FortiAP unit CLI through the FortiGate unit - CLI

1. Use the FortiGate CLI `execute telnet` command to access the FortiAP. For example, if the FortiAP unit IP address is 192.168.1.2, enter:

```
execute telnet 192.168.1.2
```
2. At the FortiAP login prompt, enter `admin`. When you are finished using the FortiAP CLI, enter `exit`.



When a WiFi controller has taken control of the FortiAP unit, Telnet access to the FortiAP unit's CLI is no longer available.

Checking and updating FortiAP unit firmware

You can view and update the FortiAP unit's firmware from the FortiGate unit that acts as its WiFi controller.

Checking the FortiAP unit firmware version

Go to **WiFi & Switch Controller > Managed FortiAPs** to view the list of FortiAP units that the FortiGate unit can manage. The **OS Version** column shows the current firmware version running on each AP.

Updating FortiAP firmware from the FortiGate unit

You can update the FortiAP firmware using either the web-based manager or the CLI. Only the CLI method can update all FortiAP units at once.

To update FortiAP unit firmware - web-based manager

1. Go to **WiFi & Switch Controller > Managed FortiAPs**.
2. Right-click the FortiAP unit in the list and select **Upgrade Firmware**.
or
Edit the FortiAP entry and select **Upgrade from File** in **FortiAP OS Version**.
3. Select **Browse** and locate the firmware upgrade file.
4. Select **OK**.
5. When the upgrade process completes, select **OK**.
The FortiAP unit restarts.

To update FortiAP unit firmware - CLI

1. Upload the FortiAP image to the FortiGate unit.

For example, the Firmware file is FAP_22A_v4.3.0_b0212_fortinet.out and the server IP address is 192.168.0.100.

```
execute wireless-controller upload-wtp-image tftp FAP_22A_v4.3.0_b0212_fortinet.out 192.168.0.100
```

If your server is FTP, change `tftp` to `ftp`, and if necessary add your user name and password at the end of the command.

2. Verify that the image is uploaded:

```
execute wireless-controller list-wtp-image
```

3. Upgrade the FortiAP units:

```
exec wireless-controller reset-wtp all
```

If you want to upgrade only one FortiAP unit, enter its serial number instead of `all`.

Updating FortiAP firmware from the FortiAP unit

You can connect to a FortiAP unit's internal CLI to update its firmware from a TFTP server on the same network. This method does not require access to the wireless controller.

1. Place the FortiAP firmware image on a TFTP server on your computer.
2. Connect the FortiAP unit to a separate private switch or hub or directly connect to your computer via a cross-over cable.
3. Change your computer's IP address to 192.168.1.3.
4. Telnet to IP address 192.168.1.2.
This IP address is overwritten if the FortiAP is connected to a DHCP environment. Ensure that the FortiAP unit is in a private network with no DHCP server.
5. Login with the username "admin" and no password.
6. Enter the following command.

For example, the FortiAP image file name is FAP_22A_v4.3.0_b0212_fortinet.out.

```
restore FAP_22A_v4.3.0_b0212_fortinet.out 192.168.1.3
```

Advanced WiFi controller discovery

A FortiAP unit can use any of six methods to locate a controller. By default, FortiAP units cycle through all six of the discovery methods. In most cases there is no need to make configuration changes on the FortiAP unit.

There are exceptions. The following section describes the WiFi controller discovery methods in more detail and provides information about configuration changes you might need to make so that discovery will work.

Controller discovery methods

There are six methods that a FortiAP unit can use to discover a WiFi controller. Below is the list of AC discovery methods used in sequence, if the FortiAP's discovery type is set to auto:

1(static) → 2(dhcp) → 3(dns) → 7(forticloud) → 5(multicast) → 6(broadcast)

For every discovery type, FortiAP sends out discovery requests and sets a timer, an interval defined as a random number of seconds (between 2-180, default is 5 seconds), which is set via the CLI:

CLI syntax

```
config wireless-controller timers
    set discovery-interval 5
end
```

After the timeout is reached, FortiAP sends out another discovery request, up to a maximum of 3 times.

After about 3 - 15 seconds, if FortiAP has no AC connection, it will switch to another discovery type and repeat the above process until the last one (**broadcast**) fails, which will lead to SULKING state.

After about 30 seconds, FortiAP will go into an AC_IP_DISCOVER state. After the AC IP is found, it will go to IDLE state, and will eventually go to the DISCOVERY state, and repeat the above process again.

Note that, while the process above is showcasing the auto discovery method, it's recommended to set the AC_DISCOVERY_TYPE to your used method in order to reduce downtime.

Static IP configuration

If FortiAP and the controller are not in the same subnet, broadcast and multicast packets cannot reach the controller. The admin can specify the controller's static IP on the AP unit. The AP unit sends a discovery request message in unicast to the controller. Routing must be properly configured in both directions.

To specify the controller's IP address on a FortiAP unit

```
cfg -a AC_IPADDR_1="192.168.0.100"
```

By default, the FortiAP unit receives its IP address, netmask, and gateway address by DHCP. If you prefer, you can assign these statically.

To assign a static IP address to the FortiAP unit

```
cfg -a ADDR_MODE=STATIC
cfg -a AP_IPADDR="192.168.0.100"
cfg -a AP_NETMASK="255.255.255.0"
cfg -a IPGW=192.168.0.1
cfg -c
```

For information about connecting to the FortiAP CLI, see [Connecting to the FortiAP CLI on page 1097](#).

DHCP

If you use DHCP to assign an IP address to your FortiAP unit, you can also provide the WiFi controller IP address at the same time. This is useful if the AP is located remotely from the WiFi controller and other discovery techniques will not work.

When you configure the DHCP server, configure Option 138 to specify the WiFi controller IP address. You need to convert the address into hexadecimal. Convert each octet value separately from left to right and concatenate them. For example, 192.168.0.1 converts to C0A80001.

If Option 138 is used for some other purpose on your network, you can use a different option number if you configure the AP units to match.

To change the FortiAP DHCP option code

To use option code 139 for example, enter

```
cfg -a AC_DISCOVERY_DHCP_OPTION_CODE=139
```

For information about connecting to the FortiAP CLI, see [Connecting to the FortiAP CLI on page 1097](#).

DNS

The access point can discover controllers through your domain name server (DNS). For the access point to do so, you must configure your DNS to return controller IP addresses in response. Allow DNS lookup of the hostname configured in the AP by using the AP parameter "AC_HOSTNAME_1".

FortiCloud

The access point can discover FortiCloud by doing a DNS lookup of the hardcoded FortiCloud AP controller hostname "apctrl1.fortinet.com". The forticloud AC discovery technique finds the AC info from apctrl1.fortinet.com using HTTPS.

FortiCloud APController: apctrl1.fortinet.com:443 208.91.113.187:443

Broadcast request

The AP unit broadcasts a discovery request message to the network and the controller replies. The AP and the controller must be in the same broadcast domain. No configuration adjustments are required.

Multicast request

The AP unit sends a multicast discovery request and the controller replies with a unicast discovery response message. The AP and the controller do not need to be in the same broadcast domain if multicast routing is properly configured.

The default multicast destination address is 224.0.1.140. It can be changed through the CLI. The address must be same on the controller and AP.

To change the multicast address on the controller

```
config wireless-controller global
  set discovery-mc-addr 224.0.1.250
end
```

To change the multicast address on a FortiAP unit

```
cfg -a AC_DISCOVERY_MC_ADDR="224.0.1.250"
```

For information about connecting to the FortiAP CLI, see [Advanced WiFi controller discovery on page 1099](#).

Wireless client load balancing for high-density deployments

Wireless load balancing allows your wireless network to distribute wireless traffic more efficiently among wireless access points and available frequency bands. FortiGate wireless controllers support the following types of client load balancing:

- Access Point Hand-off - the wireless controller signals a client to switch to another access point.
- Frequency Hand-off - the wireless controller monitors the usage of 2.4GHz and 5GHz bands, and signals clients to switch to the lesser-used frequency.

Load balancing is not applied to roaming clients.

Access point hand-off

Access point handoff wireless load balancing involves the following:

- If the load on an access point (ap1) exceeds a threshold (of for example, 30 clients) then the client with the weakest signal will be signaled by wireless controller to drop off and join another nearby access point (ap2).
- When one or more access points are overloaded (for example, more than 30 clients) and a new client attempts to join a wireless network, the wireless controller selects the least busy access point that is closest to the new client and this access point is the one that responds to the client and the one that the client joins.

Frequency hand-off or band-steering

Encouraging clients to use the 5GHz WiFi band if possible enables those clients to benefit from faster interference-free 5GHz communication. The remaining 2.4GHz clients benefit from reduced interference.

The WiFi controller probes clients to determine their WiFi band capability. It also records the RSSI (signal strength) for each client on each band.

If a new client attempts to join the network, the controller looks up that client's MAC address in its wireless device table and determines if it's a dual band device. If it is not a dual band device, then its allowed to join. If it is a dual band device, then its RSSI on 5GHz is used to determine whether the device is close enough to an access point to benefit from movement to 5GHz frequency.

If both conditions of 1) dual band device and 2) RSSI value is strong, then the wireless controller does not reply to the join request of the client. This forces the client to retry a few more times and then timeout and attempt to join the same SSID on 5GHz. Once the Controller see this new request on 5GHz, the RSSI is again measured and the client is allowed to join. If the RSSI is below threshold, then the device table is updated and the controller forces the client to timeout again. A client's second attempt to connect on 2.4GHz will be accepted.

Configuration

From the web-based manager, edit a custom AP profile and select **Frequency Handoff** and **AP Handoff** as required for each radio on the AP.

From the CLI, you configure wireless client load balancing thresholds for each custom AP profile. Enable access point hand-off and frequency hand-off separately for each radio in the custom AP profile.

```
config wireless-controller wtp-profile
  edit new-ap-profile
    set handoff-rssi <rssi_int>
    set handoff-sta-thresh <clients_int>
    config radio-1
      set frequency-handoff {disable | enable}
      set ap-handoff {disable | enable}
    end
    config radio-2
      set frequency-handoff {disable | enable}
      set ap-handoff {disable | enable}
    end
  end
end
```

Where:

- `handoff-rssi` is the RSSI threshold. Clients with a 5 GHz RSSI threshold over this value are load balanced to the 5GHz frequency band. Default is 25. Range is 20 to 30.

- `handoff-sta-thresh` is the access point handoff threshold. If the access point has more clients than this threshold it is considered busy and clients are changed to another access point. Default is 30, range is 5 to 25.
- `frequency-handoff` enable or disable frequency handoff load balancing for this radio. Disabled by default.
- `ap-handoff` enable or disable access point handoff load balancing for this radio. Disabled by default.

Frequency handoff must be enabled on the 5GHz radio to learn client capability.

FortiAP groups

FortiAP groups facilitate the application of FortiAP profiles to large numbers of FortiAPs. A FortiAP can belong to no more than one FortiAP group. A FortiAP group can include only one model of FortiAP.

Through the VLAN pool feature, a FortiAP group can be associated with a VLAN to which WiFi clients will be assigned. For more on VLAN pool assignment, see [VLAN assignment by VLAN pool](#).

FortiAP groups are only configurable in the CLI Console.

To create a FortiAP group - CLI

In this example, `wtp-group-1` is created for a FortiAP-221C and one member device is added.

```
config wireless-controller wtp-group
  edit wtp-group-1
    set platform-type 221C
    config wtp-list
      edit FP221C3X14019926
    end
  end
end
```

LAN port options

Some FortiAP models have one or more LAN interfaces that can provide wired network access. LAN ports can be

- bridged to the incoming WAN interface
- bridged to one of the WiFi SSIDs that the FortiAP unit carries
- connected by NAT to the incoming WAN interface

There are some differences among FortiAP models.

Models like 11C and 14C have one port labeled WAN and one or more ports labeled LAN. By default, the LAN ports are offline. You can configure LAN port operation in the FortiAP Profile in the GUI (**Wireless Controller > FortiAP Profiles**) or in the CLI (`config wireless-controller wtp-profile, config lan` subcommand).

Models like 320C, 320B, 112D, and 112B have two ports, labeled LAN1 and LAN2. LAN1 acts as a WAN port connecting the FortiAP to a FortiGate or FortiCloud. By default, LAN2 is bridged to LAN1. Other modes of LAN2 operation must be enabled in the CLI:

```
config wireless-controller wtp-profile
  edit <profile_name>
    set wan-port-mode wan-lan
  end
```

By default `wan-port-mode` is set to `wan-only`.

When `wan-port-mode` is set to `wan-lan`, LAN2 Port options are available in the GUI and the CLI the same as the other FortiAP models that have labeled WAN and LAN ports.

Bridging a LAN port with an SSID

Bridging a LAN port with a FortiAP SSID combines traffic from both sources to provide a single broadcast domain for wired and wireless users.

In this configuration

- The IP addresses for LAN clients come from the DHCP server that serves the wireless clients.
- Traffic from LAN clients is bridged to the SSID's VLAN. Dynamic VLAN assignment for hosts on the LAN port is not supported.
- Wireless and LAN clients are on the same network and can communicate locally, via the FortiAP.
- Any host connected to the LAN port will be taken as authenticated. RADIUS MAC authentication for hosts on the LAN port is not supported.

For configuration instructions, see [LAN port options on page 1103](#).

Bridging a LAN port with the WAN port

Bridging a LAN port with the WAN port enables the FortiAP unit to be used as a hub which is also an access point.

In this configuration

- The IP addresses for LAN clients come from the WAN directly and will typically be in the same range as the AP itself.
- All LAN client traffic is bridged directly to the WAN interface.
- Communication between wireless and LAN clients can only occur if a policy on the FortiGate unit allows it.

For configuration instructions, see [LAN port options on page 1103](#).

Configuring FortiAP LAN ports

You can configure FortiAP LAN ports for APs in a FortiAP Profile. A profile applies to APs that are the same model and share the same configuration. If you have multiple models or different configurations, you might need to create several FortiAP Profiles. For an individual AP, it is also possible to override the profile settings.

To configure FortiAP LAN ports - web-based manager

1. If your FortiAP unit has LAN ports, but no port labeled WAN (models 320C, 320B, 112D, and 112B for example), enable LAN port options in the CLI:

```
config wireless-controller wtp-profile
  edit <profile_name>
    set wan-port-mode wan-lan
  end
```
2. Go to **WiFi & Switch Controller > FortiAP Profiles**.
3. Edit the default profile for your FortiAP model or select **Create New**.
4. If you are creating a new profile, enter a **Name** and select the correct **Platform** (model).
5. Select SSIDs.
6. In the **LAN Port** section, set **Mode** to **Bridge to** and select an SSID or **WAN Port** as needed.
On some models with multiple LAN ports, you can set **Mode** to **Custom** and configure the LAN ports individually.

Enable each port that you want to use and select an SSID or **WAN Port** as needed.

7. Select OK.

Be sure to select this profile when you authorize your FortiAP units.

To configure FortiAP LAN ports - CLI

In this example, the default FortiAP-11C profile is configured to bridge the LAN port to the office SSID.

```
config wireless-controller wtp-profile
  edit FAP11C-default
    config lan
      set port-mode bridge-to-ssid
      set port-ssid office
    end
  end
end
```

In this example, the default FortiAP-28C profile is configured to bridge LAN port1 to the office SSID and to bridge the other LAN ports to the WAN port.

```
config wireless-controller wtp-profile
  edit FAP28C-default
    config lan
      set port1-mode bridge-to-ssid
      set port1-ssid office
      set port2-mode bridge-to-wan
      set port3-mode bridge-to-wan
      set port4-mode bridge-to-wan
      set port5-mode bridge-to-wan
      set port6-mode bridge-to-wan
      set port7-mode bridge-to-wan
      set port8-mode bridge-to-wan
    end
  end
end
```

In this example, the default FortiAP-320C profile is configured to bridge the LAN port to the office SSID.

```
config wireless-controller wtp-profile
  edit FAP320C-default
    set wan-port-mode wan-lan
    config lan
      set port-mode bridge-to-ssid
      set port-ssid office
    end
  end
end
```

To configure FortiAP unit LAN ports as a FortiAP Profile override - web-based manager

1. Go to **WiFi & Switch Controller > Managed FortiAPs**.
2. Select the FortiAP unit from the list and select **Edit**.
3. Select the **FortiAP Profile**, if this has not already been done.
4. In the **LAN Port** section, select **Override**.
The options for **Mode** are shown.
5. Set **Mode** to **Bridge to** and select an SSID or **WAN Port**, or **NAT to WAN** as needed.
On some models with multiple LAN ports, you can set **Mode** to **Custom** and configure the LAN ports individually.

Enable and configure each port that you want to use.

6. Select OK.

To configure FortiAP unit LAN ports as a FortiAP Profile override - CLI

In this example, a FortiAP unit's configuration overrides the FortiAP Profile to bridge the LAN port to the office SSID.

```
config wireless-controller wtp
  edit FP320C3X14020000
    set wtp-profile FAP320C-default
    set override-wan-port-mode enable
    set wan-port-mode wan-lan
    set override-lan enable
    config lan
      set port-mode bridge-to-ssid
      set port-ssid office
    end
  end
end
```

Preventing IP fragmentation of packets in CAPWAP tunnels

A common problem with controller-based WiFi networks is reduced performance due to IP fragmentation of the packets in the CAPWAP tunnel.

Fragmentation can occur because of CAPWAP tunnel overhead increasing packet size. If the original wireless client packets are close to the maximum transmission unit (MTU) size for the network (usually 1500 bytes for Ethernet networks unless jumbo frames are used) the resulting CAPWAP packets may be larger than the MTU, causing the packets to be fragmented. Fragmenting packets can result in data loss, jitter, and decreased throughput.

The FortiOS/FortiAP solution to this problem is to cause wireless clients to send smaller packets to FortiAP devices, resulting in 1500-byte CAPWAP packets and no fragmentation. The following options configure CAPWAP IP fragmentation control:

```
config wireless-controller wtp-profile
  edit FAP321C-default
    set ip-fragment-preventing {tcp-mss-adjust | icmp-unreachable}
    set tun-mtu-uplink {0 | 576 | 1500}
    set tun-mtu-downlink {0 | 576 | 1500}
  end
end
```

By default, `tcp-mss-adjust` is enabled, `icmp-unreachable` is disabled, and `tun-mtu-uplink` and `tun-mtu-downlink` are set to 0.

To set `tun-mtu-uplink` and `tun-mtu-downlink`, use the default TCP MTU value of 1500. This default configuration prevents packet fragmentation because the FortiAP unit limits the size of TCP packets received from wireless clients so the packets don't have to be fragmented before CAPWAP encapsulation.

The `tcp-mss-adjust` option causes the FortiAP unit to limit the maximum segment size (MSS) of TCP packets sent by wireless clients. The FortiAP does this by adding a reduced MSS value to the SYN packets sent by the FortiAP unit when negotiating with a wireless client to establish a session. This results in the wireless client sending packets that are smaller than the `tun-mtu-uplink` setting, so that when the CAPWAP headers are added, the CAPWAP packets have an MTU that matches the `tun-mtu-uplink` size.

The `icmp-unreachable` option affects all traffic (UDP and TCP) between wireless clients and the FortiAP unit. This option causes the FortiAP unit to drop packets that have the "Don't Fragment" bit set in their IP header and that are large enough to cause fragmentation and then send an ICMP packet -- type 3 "ICMP Destination unreachable" with code 4 "Fragmentation Needed and Don't Fragment was Set" back to the wireless controller. This should cause the wireless client to send smaller TCP and UDP packets.

Overriding IP fragmentation settings on a FortiAP

If the FortiAP Profile settings for IP fragmentation are not appropriate for a particular FortiAP, you can override the settings on that specific unit.

```
config wireless-controller wtp
  edit FAP321C3X14019926
    set override-ip-fragment enable
    set ip-fragment-preventing {tcp-mss-adjust | icmp-unreachable}
    set tun-mtu-uplink {0 | 576 | 1500}
    set tun-mtu-downlink {0 | 576 | 1500}
  end
end
```

LED options

Optionally, the status LEDs on the FortiAP can be kept dark. This is useful in dormitories, classrooms, hotels, medical clinics, hospitals where the lights might be distracting or annoying to occupants.

On the FortiGate, the LED state is controlled in the FortiAP Profile. By default the LEDs are enabled. The setting is CLI-only. For example, to disable the LEDs on FortiAP-221C units controlled by the FAP221C-default profile, enter:

```
config wireless-controller wtp-profile
  edit FAP221C-default
    set led-state disable
  end
```

You can override the FortiAP Profile LED state setting on an individual FortiAP using the CLI. For example, to make sure the LEDs are disabled on one specific unit, enter:

```
config wireless-controller wtp
  edit FAP221C3X14019926
    set override-led-state enable
    set led-state disable
  end
```

The LED state is also controllable from the FortiAP unit itself. By default, the FortiAP follows the FortiAP Profile setting.

LED schedules

Use the command below (`led-schedule`) to assign recurring firewall schedules for illuminating LEDs on the FortiAP. This entry is only available when `led-state` is enabled, at which point LEDs will be visible when at least one of the schedules is valid.

Separate multiple schedule names with a space, as configured under `config firewall schedule group` and `config firewall schedule recurring`.

Syntax

```

config wireless-controller wtp-profile
  edit {name}
    set led-state {enable | disable}
    set led-schedules <name>
  next
end

```

CAPWAP bandwidth formula

The following section provides information on how to calculate the control plane CAPWAP traffic load in local bridging. The formula provided can help estimate the approximate package bandwidth cost. This is important for knowing precisely how much bandwidth is required on a WAN link for a centralized FortiGate managing hundreds of access points.

There are multiple factors that might affect the volume of CAPWAP control traffic, including the number of stations there are and large WiFi events.

The Ethernet/IP/UDP/CAPWAP uplink header cost should be approximately 66 bytes.

The tables below depict basic and commonly used optional CAPWAP bandwidth costs, on a per-AP basis.

Note the following:

- **STA:** The number of stations associated with the FortiAP.
- **ARP scan:** Finds hidden devices in your network.
- **VAP:** The number of VAPS held by the FortiAP.
- **Radio:** The number of radios (maximum of two) enabled by the FortiAP.

Basic per-AP CAPWAP bandwidth costs

Content	Time (seconds)	Payload (byte)	Package bandwidth cost (bps)
Echo Req	30	16	$(66+16)*8/30=21.86$
STA scan	30	$25+20*sta$	$(66+25+20*sta)*8/30=24.26+5.3*sta$
ARP scan	30	$25+18*sta$	$(66+25+18*sta)*8/30=24.26+4.8*sta$
STA CAP	30	$25+19*sta$	$(66+25+19*sta)*8/30=24.26+5.1*sta$
STA stats	1	$25+41*sta$	$(66+25+41*sta)*8/1=728.0+328.0*sta$
VAP stats	15	$40+18*vap$	$(66+40+18*vap)*8/15=56.53+9.6*vap$
Radio stats	15	$25+25*radio$	$(66+25+25*radio)*8/15=48.53+13.3*radio$
Total:			$908.7+343.2*sta+9.6*vap+13.3*radio$

Commonly used optional per-AP CAPWAP bandwidth costs

Content	Time (seconds)	Payload (byte)	Package bandwidth cost (bps)
AP scan	30	25+63*scanned-ap	$(66+25+63*\text{scanned-ap}) * 8 / 30 = 24.26 + 16.8 * \text{scanned-ap}$
Total:			$932.96 + 343.2 * \text{sta} + 9.6 * \text{vap} + 13.3 * \text{radio} + 16.8 * \text{scanned-ap}$



Enabling WIDS features, LLDP, MESH, FortiPresence, and Client Station Locating Service can lead to additional bandwidth consumption.

Example:

There are 100 FortiAPs, with 187 stations distributed among them. Each FortiAP holds five VAPs among their radios, and each enables two radios. The basic CAPWAP bandwidth cost would be:

$$908.7 * 100 + 343.2 * 187 + 9.6 * 5 * 100 + 13.3 * 2 * 100 = \mathbf{162.51\text{kbps}}$$

Additionally, if two FortiAPs enabled "AP scan", and suppose one scans 99 APs in each scan and the other scans 20 APs in each scan, the additional CAPWAP bandwidth cost would be:

$$(24.26 + 16.8 * 99) + (24.26 + 16.8 * 20) = \mathbf{2\text{ kbps}}$$

Enabling LLDP protocol

You can enable the LLDP protocol in the FortiAP Profile via the CLI. Each FortiAP using that profile can then send back information about the switch and port that it is connected to.

To enable LLDP, enter the following:

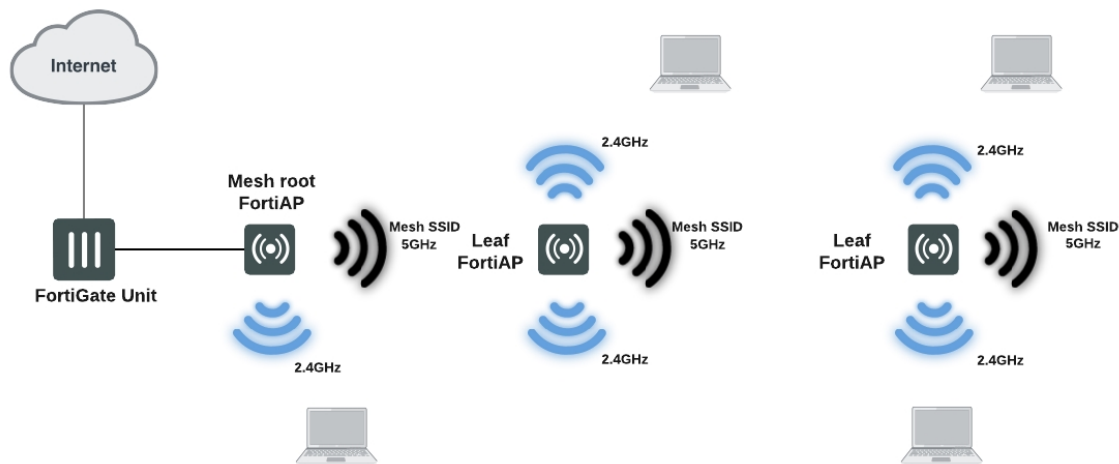
```
config wireless-controller wtp-profile
  edit <profile-name>
    set lldp enable
  end
```

Wireless mesh

The access points of a WiFi network are usually connected to the WiFi controller through Ethernet wiring. A wireless mesh eliminates the need for Ethernet wiring by connecting WiFi access points to the controller by radio. This is useful where installation of Ethernet wiring is impractical.

Overview of wireless mesh

The figure below shows a wireless mesh topology.



A wireless mesh is a multiple access point (AP) network in which only one FortiAP unit is connected to the wired network. The other FortiAPs communicate with the controller over a separate backhaul SSID that isn't available to regular WiFi clients. The AP connected to the network by Ethernet is called the mesh root node. The backhaul SSID carries CAPWAP discovery, configuration, and other communications that would usually be carried on an Ethernet connection.

The root node can be a FortiAP unit or the built-in AP of a FortiWiFi unit. APs that serve regular WiFi clients are called leaf nodes. Leaf APs also carry the mesh SSID for more distant leaf nodes. A leaf node can connect to the mesh SSID directly from the root node or from any of the other leaf nodes. This provides redundancy in case of an AP failure.

All access points in a wireless mesh configuration must have at least one of their radios configured to provide mesh backhaul communication. As with wired APs, when mesh APs start up, they can be discovered by a FortiGate or FortiWiFi unit WiFi controller and authorized to join the network.

The backhaul SSID delivers the best performance when it is carried on a dedicated radio. On a two-radio FortiAP unit, for example, the 5GHz radio could carry only the backhaul SSID while the 2.4GHz radio carries one or more SSIDs that serve users. You can configure background WiFi scanning in this mode.

The backhaul SSID can also share the same radio with SSIDs that serve users. Performance is reduced because the backhaul and user traffic compete for the available bandwidth. Background WiFi scanning isn't available in this mode. One advantage of this mode is that a two-radio AP can offer WiFi coverage on both bands.

Wireless mesh deployment modes

There are two common wireless mesh deployment modes:

Wireless mesh	Access points are connected to a FortiGate or FortiWiFi unit WiFi controller. WiFi users connect to wireless SSIDs in the same way as on non-mesh WiFi networks.
Wireless bridging	Two LAN segments are connected together over a wireless link (the backhaul SSID). On the leaf AP, the Ethernet connection can be used to provide a wired network. Both WiFi and wired users on the leaf AP are connected to the LAN segment to which the root AP is connected.


Firmware requirements

All FortiAP units that are part of the wireless mesh network must be upgraded to FortiAP firmware version 5.0, build 003, or higher. FortiAP-222B units must have their BIOS upgraded to version 400012. The FortiWiFi or FortiGate unit used as the WiFi controller must be running FortiOS firmware version 5.0 or higher.


Types of wireless mesh

A WiFi mesh can provide access to widely-distributed clients. The mesh root AP which is directly connected to the WiFi controller can be either a FortiAP unit or the built-in AP of a FortiWiFi unit that is also the WiFi controller.

FortiAP units used as both mesh root AP and leaf AP




FortiWiFi unit as mesh root AP with FortiAP units as leaf APs



An alternate use of the wireless mesh is as a point-to-point relay. Both wired and WiFi users on the leaf AP side are connected to the LAN segment on the mesh root side.

Point-to-point wireless mesh



Fast-roaming for mesh backhaul link

Mesh implementations for leaf FortiAP can perform background scanning when the leaf AP is associated with the root. Various options for background scanning can be configured with the CLI. See [Mesh variables on page 1219](#) for more details.

Configuring a meshed WiFi network

You need to:

- Create the mesh root SSID.
- Create the FortiAP profile.
- Configure mesh leaf AP units.
- Configure the mesh root AP, either a FortiWiFi unit's local radio or a FortiAP unit.
- Authorize the mesh leaf units when they connect to the WiFi Controller.
- Create security policies.

This section assumes that the end-user SSIDs already exist.

Creating the mesh root SSID

The mesh route SSID is the radio backhaul that conveys the user SSID traffic to the leaf FortiAPs.

To configure the mesh root SSID

1. Go to **WiFi & Switch Controller > SSID** and select **Create New > SSID**.
2. Enter a **Name** for the WiFi interface.
3. In **Traffic Mode**, select **Mesh Downlink**.
4. Enter the **SSID**.
5. Set **Security Mode** to **WPA2 Personal** and enter the **Pre-shared key**.
Remember the key, you need to enter it into the configurations of the leaf FortiAPs.
6. Select **OK**.

Creating the FortiAP profile

Create a FortiAP profile for the meshed FortiAPs. If more than one FortiAP model is involved, you need to create a profile for each model. Typically, the profile is configured so that Radio 1 (5GHz) carries the mesh backhaul SSID while Radio 2 (2.4GHz) carries the SSIDs to which users connect.

For Radio 1, use the **Select SSIDs** option and choose only the backhaul SSID. The radio that carries the backhaul traffic must not carry other SSIDs.

Radio 2 carries user SSIDs and shouldn't carry the backhaul. Use the **Select SSIDs** option and choose the networks that you want to provide.

For more information, see [Configuring a WiFi LAN on page 1065](#).

Configuring the mesh root AP

The mesh root AP can be either a FortiWiFi unit's built-in AP or a FortiAP unit.

To enable a FortiWiFi unit's local radio as mesh root

1. On the FortiWiFi unit, go to **WiFi & Switch Controller > Local WiFi Radio**.
2. Select **Enable WiFi Radio**.
3. In **SSID**, select **Select SSIDs**, then select the mesh root SSID.

4. Optionally, adjust **Tx Power** or select **Auto Tx Power Control**.
5. Select **Apply**.



In a network with multiple wireless controllers, make sure that each mesh root has a unique SSID. Other controllers using the same mesh root SSID might be detected as fake or rogue APs. Go to **WiFi & Switch Controller > SSID** to change the SSID.

To configure a network interface for the mesh root FortiAP unit

1. On the FortiGate unit, go to **Network > Interfaces**.
2. Edit the interface where you will connect the FortiAP unit.
3. Make sure that **Role** is **LAN**.
4. In **Addressing mode**, select **Dedicated to FortiSwitch**.
5. In **IP/Network Mask**, enter an IP address and netmask for the interface.
DHCP will provide addresses to connected devices. To maximize the number of available addresses, the interface address should end with 1, for example 192.168.10.1.
6. Select **OK**.

At this point you can connect the mesh root FortiAP (see below). If you're going to configure leaf FortiAPs through the wireless controller (see ["Configuring a meshed WiFi network" on page 1113](#)), connect the root unit later.

To enable the root FortiAP unit

1. Connect the root FortiAP unit's Ethernet port to the FortiGate network interface that you configured.
2. On the FortiGate unit, go to **WiFi & Switch Controller > Managed FortiAPs**.
If the root FortiAP unit isn't listed, wait 15 seconds and select **Refresh**. Repeat if necessary. If the unit is still missing after a minute or two, power cycle the root FortiAP unit and try again.
3. Right-click the FortiAP entry and choose your profile from the **Assign Profile** submenu.
4. Right-click the FortiAP entry and select **Authorize**.
Initially, the **State** of the FortiAP unit is **Offline**. Periodically click **Refresh** to update the status. Within about two minutes, the state changes to **Online**.
5. Select **OK**.

Configuring the mesh leaf FortiAPs

The FortiAP units that serve as leaf nodes must be preconfigured. This involves changing the FortiAP unit's internal configuration. You can do this by direct connection or through the FortiGate wireless controller.

Method 1: Direct connection to the FortiAP

1. Configure the computer's IP as 192.168.1.3.
2. Connect the computer to the FortiAP unit's Ethernet port and use the default IP address, 192.168.1.2.
3. Log in to the FortiAP as admin. By default, no password is set.
4. Enter the following commands:
 - a. If you're using the GUI, go to **Connectivity > Uplink** and select the **Mesh** option. Then enter the **Mesh AP SSID** and **Mesh AP Password** (pre-shared key).
 - b. If you're using the FortiAP CLI (Telnet or SSH), enter the following commands, substituting your own SSID and password (pre-shared key):

```
cfg -a MESH_AP_TYPE=1
cfg -a MESH_AP_SSID=fortinet.mesh.root
cfg -a MESH_AP_PASSWD=hardtoguess
cfg -c
exit
```

5. Disconnect the computer.
6. Power down the FortiAP.
7. Repeat the preceding steps for each leaf FortiAP.

Method 2: Connecting through the FortiGate unit

1. Connect the Ethernet port on the leaf FortiAP to the FortiGate network interface that you configured for FortiAPs. Connect the FortiAP unit to a power source unless PoE is used.
2. On the FortiGate unit, go to **WiFi & Switch Controller > Managed FortiAPs**.
If the FortiAP unit isn't listed, wait 15 seconds and select **Refresh**. Repeat if necessary. If the unit is still missing after a minute or two, power cycle the FortiAP unit and try again.
3. Select the discovered FortiAP unit and authorize it. Click **Refresh** every 10 seconds until the **State** indicator changes to **Online**.
4. Right-click the FortiAP and select **>_Connect to CLI**. The **CLI Console** window opens. Log in as "admin".
5. Enter the following commands, substituting your own SSID and password (pre-shared key):

```
cfg -a MESH_AP_TYPE=1
cfg -a MESH_AP_SSID=fortinet.mesh.root
cfg -a MESH_AP_PASSWD=hardtoguess
cfg -c
exit
```
6. Disconnect the FortiAP and delete it from the **Managed FortiAP** list.
7. Repeat the preceding steps for each leaf FortiAP.

Authorizing leaf APs

When the root FortiAP is connected and online, apply power to the preconfigured leaf FortiAPs. The leaf FortiAPs will connect themselves wirelessly to the WiFi Controller through the mesh network. You must authorize each unit.

1. On the FortiGate unit, go to **WiFi & Switch Controller > Managed FortiAPs**. Periodically select **Refresh** until the FortiAP unit is listed. This can take up to three minutes.
The **State** of the FortiAP unit should be **Waiting for Authorization**.
2. Right-click the FortiAP entry and choose your profile from the **Assign Profile** submenu.
3. Right-click the FortiAP entry and select **Authorize**.
Initially, the **State** of the FortiAP unit is **Offline**. Periodically click **Refresh** to update the status. Within about two minutes, the state changes to **Online**.

Creating security policies

To permit traffic to flow from the end-user WiFi network to the network interfaces for the Internet and other networks, you need to create security policies and enable NAT. [\(cross reference for WiFi policy creation\)](#)

Viewing the status of the mesh network

On the FortiGate unit, go to **WiFi & Switch Controller > Managed FortiAPs** to view the list of APs.

+ Create New ✎ Edit 🗑 Delete 🔄 Refresh				AP	Radio	Managed FortiAPs	2/32
Access Point	State	Connected Via	SSIDs	Channel	Clients	FortiAP Profile	
FP221C3X14019926	✓	🌐 192.168.2.3	Radio 1: example-staff Radio 2: fortinet.mesh.root	Radio1: 1 Radio2: 116	Radio 1: 0 Radio 2: 0	mesh-profile	
FP221C3X14023979	✓	🌐 192.168.2.2	Radio 1: example-staff Radio 2: fortinet.mesh.root	Radio1: 1 Radio2: 116	Radio 1: 0 Radio 2: 1	mesh-profile	

The **Connected Via** column lists the IP address of each FortiAP and uses icons to show whether the FortiAP is connected by Ethernet or mesh.

Ethernet	🌐
Mesh	🌐

To see how the FortiGate wireless controller connects to the FortiAP, mouse over the **Connected Via** information.

Connected Via	SSIDs
🌐 192.168.2.3	Radio 1: example-staff Radio 2: fortinet.mesh.root
🌐 192.168.2.2	Radio 1: example-staff Radio 2: fortinet.mesh.root

AP Topology

```

graph TD
    FG[This FortiGate]
    FG -.-> AP1[FP221C3X14023979  
192.168.2.2]
    AP1 -.-> AP2[FP221C3X14019926  
192.168.2.3]
  
```

Configuring a point-to-point bridge

To connect two wired network segments using a WiFi link, you can create a point-to-point bridge. The effect is the same as connecting the two network segments to the same wired switch.

You need to:

- Configure a mesh-backhaul SSID and a mesh root AP as described in ["Configuring the mesh root AP" on page 1113](#).
Note: The mesh root AP for a point-to-point bridge must be a FortiAP unit, not the internal AP of a FortiWiFi unit.
- Configure a mesh leaf FortiAP as described in ["Configuring the mesh leaf FortiAPs" on page 1114](#) and add these steps to configure the Ethernet bridge:
 - If you're using the FortiAP web-based manager, select **Ethernet Bridge**
 - If you're using the FortiAP CLI, insert the following command before the line reading `cfg -c`:
`cfg -a MESH_ETH_BRIDGE=1`
- Connect the local wired network to the Ethernet port on the mesh leaf FortiAP unit. Users are assigned IP addresses from the DHCP server on the wired network connected to the mesh root FortiAP unit.



In general, the mesh-Ethernet bridge automatically detects VLAN ID tags in data packets and allows them to pass. When necessary, you can configure VLAN IDs for permanent support in a mesh-Ethernet bridge. To do this, enter the following commands in the mesh leaf FortiAP CLI:

```
cfg -a MESH_ETH_BRIDGE_VLANS=100,200,300
cfg -c
```

Hotspot 2.0

Hotspot 2.0 Access Network Query Protocol (ANQP) is a query and response protocol that defines seamless roaming services offered by an AP. The following CLI commands are available under `config wireless-controller`, to configure Hotspot 2.0 ANQP.

Syntax

```
config wireless-controller hotspot20 anqp-3gpp-cellular
  edit {name}
    config mcc-mnc-list
      edit {id}
        set id {integer}
        set mcc {string}
        set mnc {string}
      next
    next
  end

config wireless-controller hotspot20 anqp-ip-address-type
  edit {name}
    set ipv6-address-type {option}
    set ipv4-address-type {option}
  next
end

config wireless-controller hotspot20 anqp-nai-realm
  edit {name}
    config nai-list
      edit {name}
        set encoding {enable | disable}
        set nai-realm {string}
        config eap-method
          edit {index}
            set index {integer}
            set method {option}
            config auth-param
              edit {index}
                set index {integer}
                set id {option}
                set val {option}
              next
            next
          next
        next
      next
    next
  end

config wireless-controller hotspot20 anqp-network-auth-type
  edit {name}
    set auth-type {option}
    set url {string}
  next
end
```

```
config wireless-controller hotspot20 anqp-roaming-consortium
  edit {name}
    config oi-list
      edit {index}
        set index {integer}
        set oi {string}
        set comment {string}
      next
    next
  end

config wireless-controller hotspot20 anqp-venue-name
  edit {name}
    config value-list
      edit {index}
        set index {integer}
        set lang {string}
        set value {string}
      next
    next
  end

config wireless-controller hotspot20 h2qp-conn-capability
  edit {name}
    set icmp-port {option}
    set ftp-port {option}
    set ssh-port {option}
    set http-port {option}
    set tls-port {option}
    set pptp-vpn-port {option}
    set voip-tcp-port {option}
    set voip-udp-port {option}
    set ikev2-port {option}
    set ikev2-xx-port {option}
    set esp-port {option}
  next
end

config wireless-controller hotspot20 h2qp-operator-name
  edit {name}
    config value-list
      edit {index}
        set index {integer}
        set lang {string}
        set value {string}
      next
    next
  end

config wireless-controller hotspot20 h2qp-osu-provider
  edit {name}
    config friendly-name
      edit {index}
        set index {integer}
        set lang {string}
        set friendly-name {string}
      next
    next
  end
```

```
        next
        set server-uri {string}
        set osu-method {option}
        set osu-nai {string}
        config service-description
            edit {service-id}
                set service-id {integer}
                set lang {string}
                set service-description {string}
            next
        set icon {string}
    next
end

config wireless-controller hotspot20 h2qp-wan-metric
    edit {name}
        set link-status {option}
        set symmetric-wan-link {option}
        set link-at-capacity {enable | disable}
        set uplink-speed {integer}
        set downlink-speed {integer}
        set uplink-load {integer}
        set downlink-load {integer}
        set load-measurement-duration {integer}
    next
end

config wireless-controller hotspot20 hs-profile
    edit {name}
        set access-network-type {option}
        set access-network-internet {enable | disable}
        set access-network-asra {enable | disable}
        set access-network-esr {enable | disable}
        set access-network-uesa {enable | disable}
        set venue-group {option}
        set venue-type {option}
        set hessid {mac address}
        set proxy-arp {enable | disable}
        set l2tif {enable | disable}
        set pame-bi {enable | disable}
        set anqp-domain-id {integer}
        set domain-name {string}
        set osu-ssid {string}
        set gas-comeback-delay {integer}
        set gas-fragmentation-limit {integer}
        set dgaf {enable | disable}
        set deauth-request-timeout {integer}
        set wnm-sleep-mode {enable | disable}
        set bss-transition {enable | disable}
        set venue-name {string}
        set roaming-consortium {string}
        set nai-realm {string}
        set oper-friendly-name {string}
        config osu-provider
            edit {name}
                next
        set wan-metrics {string}
```

```
        set network-auth {string}
        set 3gpp-plmn {string}
        set conn-cap {string}
        set qos-map {string}
        set ip-addr-type {string}
    next
end

config wireless-controller hotspot20 icon
    edit {name}
        config icon-list
            edit {name}
                set lang {string}
                set file {string}
                set type {option}
                set width {integer}
                set height {integer}
            next
        next
    end

config wireless-controller hotspot20 qos-map
    edit {name}
        config dscp-except
            edit {index}
                set index
                set dscp
                set up
            next
        config dscp-range
            edit {index}
                set index
                set up
                set low
                set high
            next
        next
    end
```


Combining WiFi and wired networks with a software switch

Combining WiFi and wired networks with a software switch

A WiFi network can be combined with a wired LAN so that WiFi and wired clients are on the same subnet. This is a convenient configuration for users. Note that software switches are only available if your FortiGate is in Interface mode.



Wireless Mesh features cannot be used in conjunction with this configuration because they enable the FortiAP Local Bridge option.

To create the WiFi and wired LAN configuration, you need to:

- Configure the SSID so that traffic is tunneled to the WiFi controller.
- Configure a software switch interface on the FortiGate unit with the WiFi and internal network interface as members.
- Configure Captive Portal security for the software switch interface.

To configure the SSID - web-based manager

1. Go to **WiFi & Switch Controller > SSID** and select **Create New**.
2. Enter:

Interface name	A name for the new WiFi interface, <code>homenet_if</code> for example.
Traffic Mode	Tunnel to Wireless Controller
SSID	The SSID visible to users, <code>homenet</code> for example.
Security Mode Data Encryption Preshared Key	Configure security as you would for a regular WiFi network.

3. Select **OK**.
4. Go to **WiFi & Switch Controller > Managed FortiAPs**, select the FortiAP unit for editing.
5. Authorize the FortiAP unit.
The FortiAP unit can carry regular SSIDs in addition to the Bridge SSID.

To configure the SSID - CLI

This example creates a WiFi interface “homenet_if” with SSID “homenet” using WPA-Personal security, passphrase “Fortinet1”.

```
config wireless-controller vap
  edit "homenet_if"
    set vdom "root"
    set ssid "homenet"
```

```

        set security wpa-personal
        set passphrase "Fortinet1"
    end
    config wireless-controller wtp
        edit FAP22B3U11005354
            set admin enable
            set vaps "homenet_if"
        end
    end

```

To configure the FortiGate software switch - web-based manager

1. Go to **Network > Interfaces** and select **Create New > Interface**.
2. Enter:

Interface Name	A name for the new interface, <code>homenet_nw</code> for example.
Type	Software Switch
Physical Interface Members	Add <code>homenet_if</code> and the internal network interface.
Addressing mode	Select Manual and enter an address, for example <code>172.16.96.32/255.255.255.0</code>
DHCP Server	Enable and configure an address range for clients.
Security Mode	Select Captive Portal . Add the permitted User Groups .

3. Select **OK**.

To configure the FortiGate unit - CLI

```

config system interface
    edit homenet_nw
        set ip 172.16.96.32 255.255.255.0
        set type switch
        set security-mode captive-portal
        set security-groups "Guest-group"
    end
config system interface
    edit homenet_nw
        set member "homenet_if" "internal"
    end

```

VLAN configuration

If your environment uses VLAN tagging, you assign the SSID to a specific VLAN in the CLI. For example, to assign the `homenet_if` interface to VLAN 100, enter:

```

config wireless-controller vap
    edit "homenet_if"
        set vlanid 100
    end

```

Additional configuration


The configuration described above provides communication between WiFi and wired LAN users only. To provide access to other networks, create appropriate firewall policies between the software switch and other interfaces.

FortiAP local bridging (Private cloud-managed AP)

A FortiAP unit can provide WiFi access to a LAN, even when the wireless controller is located remotely. This configuration is useful for the following situations:

- Installations where the WiFi controller is remote and most of the traffic is local or uses the local Internet gateway
- Wireless-PCI compliance with remote WiFi controller
- Telecommuting, where the FortiAP unit has the WiFi controller IP address pre-configured and broadcasts the office SSID in the user's home or hotel room. In this case, data is sent in the wireless tunnel across the Internet to the office and you should enable encryption using DTLS.

Remotely-managed FortiAP providing WiFi access to local network



On the remote FortiGate wireless controller, the WiFi SSID is created with the **Bridge with FortiAP Interface** option selected. In this mode, no IP addresses are configured. The WiFi and Ethernet interfaces on the FortiAP behave as a switch. WiFi client devices obtain IP addresses from the same DHCP server as wired devices on the LAN.



The local bridge feature cannot be used in conjunction with Wireless Mesh features.

Block-Intra-SSID Traffic is available in Bridge mode. This is useful in hotspot deployments managed by a central FortiGate, but would also be useful in cloud deployments. Previously, this was only supported in Tunnel mode.

To configure a FortiAP local bridge - web-based manager

1. Go to **WiFi & Switch Controller > SSID** and select **Create New > SSID**.
2. Enter:

Interface name	A name for the new WiFi interface.
Traffic Mode	Local bridge with FortiAP's Interface
SSID	The SSID visible to users.
Security Mode Data Encryption Preshared Key	Configure security as you would for a regular WiFi network.

3. Select **OK**.
4. Go to **WiFi & Switch Controller > Managed FortiAPs** and select the FortiAP unit for editing.
5. Authorize the FortiAP unit.
The FortiAP unit can carry regular SSIDs in addition to the Bridge SSID.

SSID configured for local bridge operation

New Interface

Interface Name

Type

Traffic Mode Local bridge with FortiAP's Int...

WiFi Settings

SSID

Security Mode

Pre-shared Key (8 - 63 characters)

Allow New WiFi Client Connections When Controller Is Down ☐

Schedule

Maximum Clients ☐

Optional VLAN ID

To configure a FortiAP local bridge - CLI

This example creates a WiFi interface “branchbridge” with SSID “LANbridge” using WPA-Personal security, passphrase “Fortinet1”.

```
config wireless-controller vap
  edit "branchbridge"
    set vdom "root"
    set ssid "LANbridge"
    set local-bridging enable
    set security wpa-personal
    set passphrase "Fortinet1"
  end
config wireless-controller wtp
  edit FAP22B3U11005354
    set admin enable
    set vaps "branchbridge"
  end
```

Note that:



- Disabling local-bridging forcefully disables local-standalone. Also, disabling either local-bridging or local-standalone forcefully disables intra-vap-privacy.
 - Enabling intra-vap-privacy forcefully disables local-standalone.
 - Enabling local-standalone forcefully enables local-bridging also.
-

Continued FortiAP operation when WiFi controller connection is down

The wireless controller, or the connection to it, might occasionally become unavailable. During such an outage, clients already associated with a bridge mode FortiAP unit continue to have access to the WiFi and wired networks. Optionally, the FortiAP unit can also continue to authenticate users if the SSID meets these conditions:

- **Traffic Mode is Local bridge with FortiAP's Interface.**
In this mode, the FortiAP unit does not send traffic back to the wireless controller.
- **Security Mode is WPA2 Personal.**
These modes do not require the user database. In WPA2 Personal authentication, all clients use the same pre-shared key which is known to the FortiAP unit.
- **Allow New WiFi Client Connections When Controller is down** is enabled.
This field is available only if the other conditions have been met.

The "LANbridge" SSID example would be configured like this in the CLI:

```
config wireless-controller vap
  edit "branchbridge"
    set vdom "root"
    set ssid "LANbridge"
    set local-bridging enable
    set security wpa-personal
    set passphrase "Fortinet1"
    set local-authentication enable
  end
```

Using bridged FortiAPs to increase scalability

The FortiGate wireless controller can support more FortiAP units in local bridge mode than in the normal mode. But this is only true if you configure some of your FortiAP units to operate in remote mode, which supports only local bridge mode SSIDs.

The Managed FortiAP page (**WiFi & Switch Controller > Managed FortiAPs**) shows at the top right the current number of Managed FortiAPs and the maximum number that can be managed, “5/64” for example. The maximum number, however, is true only if all FortiAP units operate in remote mode. For more detailed information, consult the Maximum Values Table. For each FortiGate model, there are two maximum values for managed FortiAP units: the total number of FortiAPs and the number of FortiAPs that can operate in normal mode.

To configure FortiAP units for remote mode operation

1. Create at least one SSID with **Traffic Mode** set to **Local bridge with FortiAP's Interface**.
2. Create a custom AP profile that includes *only* local bridge SSIDs.
3. Configure each managed FortiAP unit to use the custom AP profile. You also need to set the FortiAP unit's `wtp-mode` to `remote`, which is possible only in the CLI. The following example uses the CLI both to set `wtp-mode` and select the custom AP profile:

```
config wireless-controller wtp
  edit FAP22B3U11005354
    set wtp-mode remote
    set wtp-profile 220B_bridge
  end
```

Using remote WLAN FortiAPs

Remote WLAN FortiAP models enable you to provide a pre-configured WiFi access point to a remote or traveling employee. Once plugged in at home or in a hotel room, the FortiAP automatically discovers the enterprise FortiGate WiFi controller over the Internet and broadcasts the same wireless SSID used in the corporate office. Communication between the WiFi controller and the FortiAP is secure, eliminating the need for a VPN.

Split tunneling

By default, all traffic from the remote FortiAP is sent to the FortiGate WiFi controller. If split tunneling is configured, only traffic destined for the corporate office networks is routed to the FortiGate. Other general Internet traffic is routed unencrypted through the local gateway. Split tunneling avoids loading the FortiGate with unnecessary traffic and allows direct access to local private networks at the location of the FortiAP even if the connection to the WiFi controller goes down.

By default, split tunneling options are not visible in the FortiGate GUI. You can make these options visible using the following CLI command:

```
config system settings
    set gui-fortiap-split-tunneling enable
end
```

Split tunneling is configured in **Managed FortiAPs**, **FortiAP Profiles**, and enabled in the **SSID**.

Configuring the FortiGate for remote FortiAPs

This section assumes that you have already defined SSIDs and now want to make them available to remote FortiAPs.

- Create FortiAP profiles for the Remote LAN FortiAP models
- If split tunneling will be used
 - configure override split tunneling in Managed FortiAPs
 - enable Split Tunneling in the SSID
 - configure the split tunnel networks in the FortiAP profile

Override Split tunneling

Go to **WiFi & Switch Controller > Managed FortiAPs** and edit your managed APs. When preconfiguring the AP to connect to your FortiGate WiFi controller, you can choose to override split tunneling, optionally including the local subnet of the FortiAP.

Creating FortiAP profiles

If you were not already using Remote LAN FortiAP models, you will need to create FortiAP profiles for them. In the FortiAP profile, you specify the SSIDs that the FortiAP will broadcast. For more information, see ["Creating a FortiAP profile" on page 1068](#).

Configuring split tunneling - FortiGate GUI

Go to **WiFi & Switch Controller > SSID** and edit your SSID. In the **WiFi Settings** section, enable **Split Tunneling**.

Go to **WiFi Controller > FortiAP Profiles** and edit the FortiAP Profile(s) that apply to the AP types used in the WiFi network. In the **Split Tunneling** section, enable **Include Local Subnet** and **Split Tunneling Subnet(s)**, where you can enter a list all of the destination IP address ranges that should **not** be routed through the the FortiGate WiFi controller. Packets for these destinations will instead be routed through the remote gateway local to the FortiAP.

The list of split tunneling subnets includes public Internet destinations and private subnets local to the FortiAP. Split tunneling public Internet destinations reduces traffic through the FortiGate unit. Split tunneling local private subnets allows these networks to be accessible to the client behind the FortiAP. Otherwise, private network IP destinations are assumed to be behind the FortiGate WiFi controller.

Configuring split tunneling - FortiGate CLI

In this example, split tunneling is configured on the example-ssid WiFi network. On FortiAP model 21D, traffic destined for the 192.168.x.x range will not be routed through the FortiGate WiFi controller. This private IP address range is typically used as a LAN by home routers.

```
config wireless-controller vap
  edit example-ssid
    set split-tunneling enable
  end

config wireless-controller wtp-profile
  edit FAP21D-default
    set split-tunneling-acl-local-ap-subnet enable
    config split-tunneling-acl
      edit 1
        set dest-ip 192.168.0.0 255.255.0.0
      end
    end
  end
```

To enter multiple subnets, create a split-tunneling-acl entry for each one.

Overriding the split tunneling settings on a FortiAP

If the FortiAP Profile split tunneling settings are not appropriate for a particular FortiAP, you can override the settings on that unit.

```
config wireless-controller wtp
  edit FAP321C3X14019926
    set override-split-tunnel enable
    set split-tunneling-acl-local-ap-subnet enable
    config split-tunneling-acl
      edit 1
        set dest-ip 192.168.10.0 255.255.255.0
      end
    end
  end
```

Configuring the FortiAP units

Prior to providing a Remote WLAN FortiAP unit to an employee, you need to preconfigure the AP to connect to your FortiGate WiFi controller.

To pre-configure a FortiAP

1. Connect the FortiAP to the FortiGate unit.
2. Go to **WiFi & Switch Controller > Managed FortiAPs** and wait for the FortiAP to be listed. Click **Refresh** periodically to see the latest information. Note the **Connected Via** IP address.
3. Go to **Dashboard**. In the CLI Console, log into the FortiAP CLI.
For example, if the IP address is 192.168.1.4, enter:

```
exec telnet 192.168.1.4
```

Enter *admin* at the login prompt. By default, no password is set.
4. Enter the following commands to set the FortiGate WiFi controller IP address. This should be the FortiGate Internet-facing IP address, in this example 172.20.120.142.

```
cfg -a AC_IPADDR_1=172.20.120.142  
cfg -c
```
5. Enter *exit* to log out of the FortiAP CLI.

Preauthorizing FortiAP units

By preauthorizing FortiAP units, you facilitate their automatic authorization on the network. Also, you can assign each unit a unique name, such as the employee's name, for easier tracking.

1. Go to **WiFi & Switch Controller > Managed FortiAPs** and create a new entry.
2. Enter the **Serial Number** of the FortiAP unit and give it a **Name**. Select the appropriate **FortiAP Profile**.
3. Click **OK**.

Repeat this process for each FortiAP.

Features for high-density deployments

High-density environments such as auditoriums, classrooms, and meeting rooms present a challenge to WiFi providers. When a large number of mobile devices try to connect to a WiFi network, difficulties arise because of the limited number of radio channels and interference between devices.

FortiOS and FortiAP devices provide several tools to mitigate the difficulties of high-density environments.

Multiple FortiAP firmware upgrades at once

Administrators can configure multiple FortiAP and FortiSwitch firmware upgrades to occur in one click (under **WiFi & Switch Controller > Managed FortiAPs**), removing the need to upgrade each device one at a time.

Power save feature

Occasionally, voice calls can become disrupted. One way to alleviate this issue is by controlling the power save feature, or to disable it altogether.

Manually configure packet transmit optimization settings by entering the following command:

```
config wireless-controller wtp-profile
edit <name>
config <radio-1> | <radio-2>
set transmit-optimize {disable | power-save | aggr-limit | retry-limit | sendbar}
```

- **disable:** Disable transmit optimization.
- **power-save:** Mark a client as power save mode if excessive transmit retries happen.
- **aggr-limit:** Set aggregation limit to a lower value when data rate is low.
- **retry-limit:** Set software retry limit to a lower value when data rate is low.
- **send-bar:** Do not send BAR frame too often.

11n radio powersave optimization

The following `powersave-optimize` parameters (under `config radio`) are used for 11n radios to optimize system performance for specific situations.

- **tim:** Set traffic indication map (TIM) bit for client in power save mode. TIM bit mask indicates to any sleeping listening stations if the AP has any buffered frames present. If enabled, the AP will always indicate to the connected client that there is a packet waiting in the AP, so it will help to prevent the client from entering a sleep state.
- **ac-vo:** Use Access Category (AC) Voice (VO) priority to send packets in the power save queue. AC VO is one of the highest classes/priority levels used to ensure quality of service (QoS). If enabled, when a client returns from a sleep state, the AP will send its buffered packet using a higher priority queue, instead of the normal priority queue.
- **no-obss-scan:** Do not put Overlapping Basic Service Set (OBSS), or high-noise (i.e. non-802.11), scan IE into a Beacon or Probe Response frame.
- **no-11b-rate:** Do not send frame using 11b data rate.
- **client-rate-follow:** Adapt transmitting PHY rate with receiving PHY rate from client. If enabled, the AP will integrate the current client's transmission PHY rate into its rate adaptation algorithm for transmitting.

Broadcast packet suppression

You can use broadcast packet suppression to reduce the traffic on your WiFi networks. In addition, some broadcast packets are unnecessary or even potentially detrimental to the network and should be suppressed. To configure broadcast suppression for each virtual access point, enter the following commands:

```
config wireless-controller vap
  edit <name>
    set broadcast-suppression {dhcp-up | dhcp-down | dhcp-starvation | arp-known | arp-unknown | arp-reply | arp-poison | arp-proxy | netbios-ns | netbios-ds | ipv6 | all-other-mc | all-other-bc}
  end
```

By default, both the `dhcp-up` and `arp-known` options are enabled. The following example leaves the default settings in place and configures a virtual access point to suppress:

- unnecessary DHCP downlink broadcast packets
- broadcast ARP requests for unknown wireless clients
- broadcast packets not covered by any of the other options

```
config wireless-controller vap
  edit <name>
    set broadcast-suppression dhcp-down arp-unknown all-other-bc
  end
```

Option	Description
<code>dhcp-up</code>	Suppress unnecessary DHCP uplink broadcast packets and prevent malicious WiFi clients from acting as DHCP servers. Default setting.
<code>dhcp-down</code>	Suppress unnecessary DHCP downlink broadcast packets and prevent malicious WiFi clients from acting as DHCP servers.
<code>dhcp-starvation</code>	Suppress DHCP starvation request messages. Prevent clients from making multiple requests and depleting the DHCP address pool.
<code>arp-known</code>	Suppress broadcast ARP for known wireless clients. Forward ARP packets to the destination client as a unicast packet and prevents them from reaching other clients. Default setting.
<code>arp-unknown</code>	Suppress broadcast ARP for unknown wireless clients.
<code>arp-reply</code>	Convert broadcast packet ARP replies to unicast and send them to clients listed on the destination MAC.
<code>arp-poison</code>	Suppress ARP poison messages from wireless clients. Prevent clients from spoofing ARP messages.
<code>arp-proxy</code>	Reply to ARP requests for wireless clients as a proxy. Generate ARP reply packets as a proxy and suppress broadcast packets. The <code>arp-known</code> option must be set for <code>arp-proxy</code> to work.
<code>netbios-ns</code>	Suppress NetBIOS name services packets with UDP port 137.

Option	Description
netbios-ds	Suppress NetBIOS datagram services packets with UDP port 138.
ipv6	Suppress IPv6 broadcast packets.
all-other-mc	Suppress multicast packets not covered by any of the specific options.
all-other-bc	Suppress broadcast packets not covered by any of the specific options.

Multicast to unicast conversion

FortiOS provides a multicast enhancement option (disabled by default) that converts multicast streams to unicast and improves performance in WiFi networks. Multicast data, such as streaming audio or video, is sent at a low data rate in WiFi networks. A unicast stream is sent to each client at high data rate that makes more efficient use of air time. To enable multicast-to-unicast conversion, enter the following commands:

```
config wireless-controller vap
  edit <vap_name>
    set multicast-enhance enable
  end
```

Ignore weak or distant clients

Clients beyond the intended coverage area can have some impact on your high-density network. Your APs will respond to these clients' probe signals, consuming valuable air time. You can configure your WiFi network to ignore weak signals that most likely come from beyond the intended coverage area. The settings are available in the CLI:

```
config wireless-controller vap
  edit <vap_name>
    set probe-resp-suppression enable
    set probe-resp-threshold <level_int>
  end
```

vap_name is the SSID name.

probe-resp-threshold is the signal strength in dBm below which the client is ignored. The range is -95 to -20dBm. The default level is -80dBm.

Turn off the 802.11b protocol

By disabling support for the obsolete 802.11b protocol, you can reduce the air time that data frames occupy. These signals will now be sent at a minimum of 6Mbps, instead of 1Mbps. You can set this for each radio in the FortiAP profile, using the CLI:

```
config wireless-controller wtp-profile
  edit <name_string>
    config radio-1
      set powersave-optimize no-11b-rate
    end
```

Disable low data rates

Each of the 802.11 protocols supports several data rates. By disabling the lowest rates, air time is conserved, allowing the channel to serve more users. You can set the available rates for each 802.11 protocol: a, b, g, n, ac. Data rates set as Basic are mandatory for clients to support. Other specified rates are supported.

The 802.11 a, b, and g protocols are specified by data rate. 802.11a can support 6, 9, 12, 18, 24, 36, 48, and 54 Mb/s. 802.11b/g can support 1, 2, 5.5, 6, 9, 12, 18, 24, 36, 48, 54 Mb/s. Basic rates are specified with the suffix "basic", "12-basic" for example. The capabilities of expected client devices need to be considered when deciding the lowest Basic rate.

The 802.11n and ac protocols are specified by the Modulation and Coding Scheme (MCS) Index and the number of spatial streams.

- 802.11n with 1 or 2 spatial streams can support mcs0/1, mcs1/1, mcs2/1, mcs3/1, mcs4/1, mcs5/1, mcs6/1, mcs7/1, mcs8/2, mcs9/2, mcs10/2, mcs11/2, mcs12/2, mcs13/2, mcs14/2, mcs15/2.
- 802.11n with 3 or 4 spatial streams can support mcs16/3, mcs17/3, mcs18/3, mcs19/3, mcs20/3, mcs21/3, mcs22/3, mcs23/3, mcs24/4, mcs25/4, mcs26/4, mcs27/4, mcs28/4, mcs29/4, mcs30/4, mcs31/4.
- 802.11ac with 1 or 2 spatial streams can support mcs0/1, mcs1/1, mcs2/1, mcs3/1, mcs4/1, mcs5/1, mcs6/1, mcs7/1, mcs8/1, mcs9/1, mcs0/2, mcs1/2, mcs2/2, mcs3/2, mcs4/2, mcs5/2, mcs6/2, mcs7/2, mcs8/2, mcs9/2.
- 802.11ac with 3 or 4 spatial streams can support mcs0/3, mcs1/3, mcs2/3, mcs3/3, mcs4/3, mcs5/3, mcs6/3, mcs7/3, mcs8/3, mcs9/3, mcs0/4, mcs1/4, mcs2/4, mcs3/4, mcs4/4, mcs5/4, mcs6/4, mcs7/4, mcs8/4, mcs9/4

Here are some examples of setting basic and supported rates.

```
config wireless-controller vap
edit <vap_name>
    set rates-11a 12-basic 18 24 36 48 54
    set rates-11bg 12-basic 18 24 36 48 54
    set rates-11n-ss34 mcs16/3 mcs18/3 mcs20/3 mcs21/3 mcs22/3 mcs23/3 mcs24/4 mcs25/4
    set rates-11ac-ss34 mcs0/3 mcs1/3 mcs2/3 mcs9/4 mcs9/3
end
```

Limit power

High-density deployments usually cover a small area that has many clients. Maximum AP signal power is usually not required. Reducing the power reduces interference between APs. Fortinet recommends that you use FortiAP automatic power control. You can set this in the FortiAP profile.

1. Go to **WiFi & Switch Controller > FortiAP Profiles** and edit the profile for your AP model.
2. For each radio, enable **Auto TX Power Control** and set the **TX Power Low** and **TX Power High** levels. The default range of 10 to 17dBm is recommended.

Use frequency band load-balancing

In a high-density environment is important to make the best use of the two WiFi bands, 2.4GHz and 5GHz. The 5GHz band has more non-overlapping channels and receives less interference from non-WiFi devices, but not all devices support it. Clients that are capable of 5GHz operation should be encouraged to use 5GHz rather than the 2.4GHz band.

To load-balance the WiFi bands, you enable Frequency Handoff in the FortiAP profile. In the FortiGate web-based manager, go to **WiFi & Switch Controller > FortiAP Profiles** and edit the relevant profile. Or, you can use the CLI:

```
config wireless-controller wtp-profile
edit FAP221C-default
config radio-1
set frequency-handoff enable
end
```

The FortiGate wireless controller continuously performs a scan of all clients in the area and records their signal strength (RSSI) on each band. When Frequency Handoff is enabled, the AP does not reply to clients on the 2.4GHz band that have sufficient signal strength on the 5GHz band. These clients can associate only on the 5GHz band. Devices that support only 2.4GHz receive replies and associate with the AP on the 2.4GHz band.

Setting the handoff RSSI threshold

The FortiAP applies load balancing to a client only if the client has a sufficient signal level on 5GHz. The minimum signal strength threshold is set in the FortiAP profile, but is accessible only through the CLI:

```
config wireless-controller wtp-profile
edit FAP221C-default
set handoff-rssi 25
end
```

`handoff-rssi` has a range of 20 to 30. RSSI is a relative measure. The higher the number, the stronger the signal.

AP load balancing

The performance of an AP is degraded if it attempts to serve too many clients. In high-density environments, multiple access points are deployed with some overlap in their coverage areas. The WiFi controller can manage the association of new clients with APs to prevent overloading.

To load-balance between APs, enable AP Handoff in the FortiAP profile. In the FortiGate web-based manager, go to **WiFi & Switch Controller > FortiAP Profiles** and edit the relevant profile. Or, you can use the CLI:

```
config wireless-controller wtp-profile
edit FAP221C-default
config radio-1
set ap-handoff enable
end
```

When an AP exceeds the threshold (the default is 30 clients), the overloaded AP does not reply to a new client that has a sufficient signal at another AP.

Setting the AP load balance threshold

The thresholds for AP handoff are set in the FortiAP profile, but is accessible only through the CLI:

```
config wireless-controller wtp-profile
edit FAP221C-default
set handoff-sta-thresh 30
set handoff-rssi 25
end
```

`handoff-sta-thresh` sets the number of clients at which AP load balancing begins. It has a range of 5 to 35.

`handoff-rssi` Sets the minimum signal strength that a new client must have at an alternate AP for the overloaded AP to ignore the client. It has a range of 20 to 30. RSSI is a relative measure. The higher the number, the stronger the signal.

Application rate-limiting

To prevent particular application types from consuming too much bandwidth, you can use the FortiOS Application Control feature.

1. Go to **Security Profiles > Application Control**.
You can use the default profile or create a new one.
2. Click the category, select **Traffic Shaping** and then select the priority for the category.
Repeat for each category to be controlled.
3. Select **Apply**.
4. Go to **Policy & Objects > IPv4 Policy** and edit your WiFi security policy.
5. In **Security Profiles**, set **Application Control** ON and select the security profile that you edited.
6. Select **OK**.

AP group management and dynamic VLAN assignment

The FortiGate can create FortiAP Groups, under **WiFi & Switch Controller > Managed FortiAPs** by selecting **Create New > Managed AP Group**, where multiple APs can be managed. AP grouping allows specific profile settings to be applied to many APs all at once that belong to a certain AP group, simplifying the administrative workload.

Note that each AP can only belong to one group.

In addition, VLANs can be assigned dynamically based on the group which an AP belongs. When defining an SSID, under **WiFi & Switch Controller > SSID**, a setting called **VLAN Pooling** can be enabled where you can either assign the VLAN ID of the AP group the device is connected to, to each device as it is detected, or to always assign the same VLAN ID to a specific device. Dynamic VLAN assignment allows the same SSID to be deployed to many APs, avoiding the need to produce multiple SSIDs.

Sharing tunnel SSIDs within a single managed AP between VDOMs as a virtual AP for multi-tenancy

This feature provides the ability to move a tunnel mode VAP into a VDOM, similar to an interface/VLAN in VDOMs. FortiAP is registered into the root VDOM.

Within a customer VDOM, customer VAPs can be created/added. In the root VDOM, the customer VAP can be added to the registered FortiAP. Any necessary firewall rules and interfaces can be configured between the two VDOMs.

Syntax

```
config wireless-controller global
  set wtp-share {enable | disable}
end
```

Manual quarantine of devices on FortiAP (tunnel mode)

Quarantined MAC addresses are blocked on the connected FortiAP from the network and the LAN. When a tunnel VAP is created, a sub-interface named **wqtn** is automatically created under tunnel interface. This sub-interface is added under a software switch.

To quarantine an SSID, go to **WiFi & Switch Controller > SSID**. Edit the SSID, and enable **Quarantine Host** is enabled under **WiFi Settings**.

Alternatively, this can be configured in the CLI Console. This feature consolidates previous CLI syntax for quarantining a host, so that the host does not need to be configured in multiple places (FortiAP and FortiSwitch). Host endpoints can be entered in a single place and the host will be quarantined throughout the access layer devices on the Fortinet Security Fabric.



Note that you can only an SSID in Tunnel Mode.

Syntax - SSID:

```
config wireless-controller vap
  edit <name>
    set quarantine {enable | disable}
  next
end
```

Syntax - Software Switch, DHCP, and User Quarantine

```
config system switch-interface
  edit "wqt.root"
    set vdom "root"
    set member "wqtn.26.AV-Qtn"
  next
end

config system dhcp server
  edit <id>
    set interface "AV-Qtn"
    config ip-range
      edit <id>
        set start-ip 10.111.0.2
        set end-ip 10.111.0.254
      next
    ...
  next
end

config user quarantine
  set quarantine {enable | disable}
end
```

To list stations in quarantine, use the following diagnose command:

```
diagnose wireless-controller wlac -c sta-qtn
```

Host quarantine per SSID

Upon creating or editing an SSID, a **Quarantine Host** option is available to enable (by default) or disable quarantining devices that are connected in Tunnel-mode. The option to quarantine a device is available on **Topology** and **FortiView** WiFi pages.

When a host is put into quarantine VLAN, it will get its IP from the quarantine VLAN's DHCP server, and become part of the quarantined network.

Syntax

```
config wireless-controller vap
  edit <name>
    set quarantine {enable | disable}
  next
end
```

To list all stations in quarantine:

```
diagnose wireless-controller wlac -c sta-qtn
```

Locate a FortiAP with LED blinking

If you have an environment that contains numerous APs, and there is one AP that you need to frequently monitor, you can configure it to blink in the FortiCloud web portal. The blinking AP will be easier to locate.

To start or stop LED blinking of a managed FortiAP, using the GUI:

1. Go to **WiFi & Switch Controller > Managed FortiAPs**.
2. Right-click in the row of the device you want to control.
3. In the dialog box, scroll down to **LED Blink** and select **Start** or **Stop**.

The following models support LED blink control through the GUI, operating on FortiAP software 6.0.1, or later:

- FortiAP-112D, 221C, 223C, 224D, 320C, 321C
- FortiAP-S/W2

To start or stop LED blinking of a managed FortiAP, using the CLI:

```
execute wireless-controller led-blink <wtp-id> {on | on 10 | off}
```

The following models support LED blink control through the CLI, operating on FortiAP software 5.6.2, or later:

- FortiAP-112D, 221C, 223C, 224D, 320C, 321C
- FortiAP-S/W2

Wireless controller optimization for large deployment - AP image upgrade

Using the CLI to upgrade FortiAP image is the preferred method especially for large deployments. Use the following execute command to upload the desired FortiAP image on the controller:

```
execute wireless-controller upload-wtp-image
```

After entering the command, reboot the FortiAP devices. This feature allows the administrator to configure all FortiAP devices to download the image from the controller at join time.

Syntax

```
config wireless-controller global
  set image-download {enable | disable}
end
```

To fine-tune this process, in order to deploy FortiAP image upgrades to a subset of devices for pilot testing, use the following command:

```
config wireless-controller wtp
```

```
edit <name>
    set image-download {enable | disable}
next
end
```

Control message off-loading and aeroscout enhancement

Users can configure control message off-loading to optimize performance. This is especially useful in environments where the AP count is around 300-350 (with a device count between 1500 and 3000), where existing users are disconnected and unable to reauthenticate due to high CPU usage. This feature includes aeroscout enhancements.

Syntax

```
config wireless-controller global
    set control-message-offload {evp-frame | areoscout-tag | ap-list | sta-list | sta-cap-
        list | stats | aeroscout-mu}
end

config wireless-controller wtp-profile
    edit <name>
        set control-message-offload {enable | disable}
        config lbs
            set ekahau-blink-mode {enable | disable}
            set aeroscout {enable | disable}
            set aeroscout-server-ip <address>
            set aeroscout-server-port <UDP listening port>
            set aeroscout-mu {enable | disable}
        end
    end
end
```

Protecting the WiFi network

Wireless IDS

The FortiGate Wireless Intrusion Detection System (WIDS) monitors wireless traffic for a wide range of security threats by detecting and reporting on possible intrusion attempts. When an attack is detected the FortiGate unit records a log message.

You can create a WIDS profile to enable these types of intrusion detection:

- Asleep Attack—ASLEAP is a tool used to perform attacks against LEAP authentication.
- Association Frame Flooding—A Denial of Service attack using a large number of association requests. The default detection threshold is 30 requests in 10 seconds.
- Authentication Frame Flooding—A Denial of Service attack using a large number of association requests. The default detection threshold is 30 requests in 10 seconds.
- Broadcasting De-authentication—This is a type of Denial of Service attack. A flood of spoofed de-authentication frames forces wireless clients to de-authenticate, then re-authenticate with their AP.
- EAPOL Packet Flooding—Extensible Authentication Protocol over LAN (EAPOL) packets are used in WPA and WPA2 authentication. Flooding the AP with these packets can be a denial of service attack. Several types of EAPOL packets are detected: EAPOL-FAIL, EAPOL-LOGOFF, EAPOL-START, EAPOL-SUCC.
- Invalid MAC OUI—Some attackers use randomly-generated MAC addresses. The first three bytes of the MAC address are the Organizationally Unique Identifier (OUI), administered by IEEE. Invalid OUIs are logged.
- Long Duration Attack—To share radio bandwidth, WiFi devices reserve channels for brief periods of time. Excessively long reservation periods can be used as a denial of service attack. You can set a threshold between 1000 and 32 767 microseconds. The default is 8200.
- Null SSID Probe Response—When a wireless client sends out a probe request, the attacker sends a response with a null SSID. This causes many wireless cards and devices to stop responding.
- Spoofed De-authentication—Spoofed de-authentication frames are a denial of service attack. They cause all clients to disconnect from the AP.
- Weak WEP IV Detection—A primary means of cracking WEP keys is by capturing 802.11 frames over an extended period of time and searching for patterns of WEP initialization vectors (IVs) that are known to be weak. WIDS detects known weak WEP IVs in on-air traffic.
- Wireless Bridge—WiFi frames with both the fromDS and ToDS fields set indicate a wireless bridge. This will also detect a wireless bridge that you intentionally configured in your network.

You can enable wireless IDS by selecting a WIDS Profile in your FortiAP profile.

To create a WIDS Profile

1. Go to **WiFi & Switch Controller > WIDS Profiles**.
2. Select a profile to edit or select **Create New**.
3. Select the types of intrusion to protect against.
By default, all types are selected.
4. Select **Apply**.

You can also configure a WIDS profile in the CLI using the `config wireless-controller wids-profile` command.

Rogue AP detection

The WIDS profile includes settings for detection of unauthorized (rogue) access points in your wireless network. For more information, see [Wireless network monitoring on page 1145](#).

WIDS client deauthentication rate for DoS attacks

As part of mitigating a Denial of Service (DoS) attack, the FortiGate sends deauthentication packets to unknown clients. In an aggressive attack, this deauthentication activity can prevent the processing of packets from valid clients. A WIDS Profile option in the CLI limits the deauthentication rate.

```
config wireless-controller wids-profile
  edit default
    set deauth-unknown-src-thresh <1-65535>
  end
```

The value set is a measure of the number of deauthorizations per second. 0 means no limit. The default is 10.

WiFi data channel encryption

Optionally, you can apply DTLS encryption to the data channel between the wireless controller and FortiAP units. This enhances security.

There are data channel encryption settings on both the FortiGate unit and the FortiAP units. At both ends, you can enable Clear Text, DTLS encryption, or both. The settings must agree or the FortiAP unit will not be able to join the WiFi network. By default, both Clear Text and DTLS-encrypted communication are enabled on the FortiAP unit, allowing the FortiGate setting to determine whether data channel encryption is used. If the FortiGate unit also enables both Clear Text and DTLS, Clear Text is used.

Data channel encryption settings are located in the FortiAP profile. By default, only Clear Text is supported.



Data channel encryption is software-based and can affect performance. Verify that the system meets your performance requirements with encryption enabled.

Configuring encryption on the FortiGate unit

You can use the CLI to configure data channel encryption.

Enabling encryption

In the CLI, the `wireless wtp-profile` command contains a new field, `dtls-policy`, with options `clear-text` and `dtls-enabled`. To enable encryption in profile1 for example, enter:

```
config wireless-controller wtp-profile
  edit profile1
    set dtls-policy dtls-enabled
  end
```

Configuring encryption on the FortiAP unit

The FortiAP unit has its own settings for data channel encryption.

Enabling CAPWAP encryption - FortiAP web-based manager

1. On the **System Information** page, in **WTP Configuration > AC Data Channel Security**, select one of:
 - Clear Text
 - DTLS Enabled
 - Clear Text or DTLS Enabled (default)
2. Select **Apply**.

Enabling encryption - FortiAP CLI

You can set the data channel encryption using the `AP_DATA_CHAN_SEC` variable: 'clear', or 'ipsec', or 'dtls'.

For example, to set security to DTLS and then save the setting, enter:

```
cfg -a AP_DATA_CHAN_SEC=dtls
cfg -c
```

Protected Management Frames and Opportunistic Key Caching support

Protected Management Frames (PMF) protect some types of management frames like deauthorization, disassociation and action frames. This feature, now mandatory on WiFi certified 802.11ac devices, prevents attackers from sending plain deauthorization/disassociation frames to disrupt or tear down a connection/association. PMF is a Wi-Fi Alliance specification based on IEEE 802.11w.

To facilitate faster roaming client roaming, you can enable Opportunistic Key Caching (OKC) on your WiFi network. When a client associates with an AP, its PMK identifier is sent to all other APs on the network. This eliminates the need for an already-authenticated client to repeat the full EAP exchange process when it roams to another AP on the same network.

Use of PMF and OKC on an SSID is configurable only in the CLI:

```
config wireless-controller vap
  edit <vap_name>
    set pmf {disable | enable | optional}
    set pmf-assoc-comeback-timeout <integer>
    set pmf-sa-query-retry-timeout <integer>
    set okc {disable | enable}
  next
end
```

When `pmf` is set to `optional`, it is considered enabled, but will allow clients that do not use PMF. When `pmf` is set to `enable`, PMF is required by all clients.

Bluetooth Low Energy (BLE) Scan

The FortiGate can configure FortiAP Bluetooth Low Energy (BLE) scan, incorporating Google's BLE beacon profile known as Eddystone, used to identify groups of devices and individual devices.



Currently, only the FAP-S221E, FAP-S223E, and FAP-222E models support this feature.

Use the following syntax to configure BLE profiles, configure BLE report intervals, and assign BLE profiles to WTP profiles.

CLI syntax - Configure BLE profiles

```
config wireless-controller ble-profile
edit <name>
    set comment <comment>
    set advertising {ibeacon | eddystone-uid | eddystone-url}
    set ibeacon-uuid <uuid>
    set major-id <0 - 65535> - (default = 1000)
    set minor-id <0 - 65535> - (default = 1000)
    set eddystone-namespace <10-byte namespace>
    set eddystone-instance <device id>
    set eddystone-url <url>
    set txpower <0 - 12> - (default = 0)
    set beacon-interval <40 - 3500> - (default = 100)
    set ble-scanning {enable | disable} - (default = disable)
next
end
```

Note that `txpower` determines the transmit power level on a scale of 0-12:

0: -21 dBm	1: -18 dBm	2: -15 dBm	3: -12 dBm	4: -9 dBm
5: -6 dBm	6: -3 dBm	7: 0 dBm	8: 1 dBm	9: 2 dBm
10: 3 dBm	11: 4 dBm	12: 5 dBm		

CLI syntax - Configure BLE report intervals

```
config wireless-controller timers
    set ble-scan-report-intv - (default = 30 sec)
end
```

CLI syntax - Assign BLE profiles to WTP profiles

```
config wireless-controller wtp-profile
edit <name>
    set ble-profile <name>
next
end
```

Preventing local bridge traffic from reaching the LAN

The following command can be enabled so that when a client connects to a VAP, and its traffic is not tunneled to the controller, the admin can control whether the client can access the local network.

Note that this entry is only available when `local-standalone-nat` is set to enable.

Syntax:

```
config wireless-controller vap
edit <name>
    set local-lan {allow | deny}
```

```

    next
end

```

FortiAP-S bridge mode security profiles

If you have enabled bridge mode for a managed FortiAP-S, you can add a UTM profile to the wireless controller configuration that allows you to apply the following security profile features to all traffic accepted by the managed FortiAP-S:

- AntiVirus (including Botnet protection),
- IPS,
- Application control, and
- Web Filtering.

You can use the following CLI command to add a wireless controller UTM profile:

```

config wireless-controller utm-profile
  edit <name>
    set comment "Default configuration for offloading WiFi traffic."
    set ips-sensor "wifi-default"
    set application-list "wifi-default"
    set antivirus-profile "wifi-default"
    set webfilter-profile "wifi-default"
    set firewall-profile-protocol-options "wifi-default"
    set firewall-ssl-ssh-profile "wifi-default"
  next
end

```

You can use the following CLI command to add a wireless controller UTM profile to a wireless SSID (virtual access point) configuration:

```

config wireless-controller vap
  edit <name>
    set utm-profile
  end
end

```

DHCP snooping and option 82 (circuit-id) options for wireless access points

New commands are available to enable or disable (by default) DHCP 82 option insertion for wireless access points. DHCP snooping is used to prevent rogue DHCP servers from offering IP addresses to DHCP clients.

Syntax

```

config wireless-controller vap
  edit wifi
    set dhcp-option82-insertion {enable | disable}
    set dhcp-option82-circuit-id-insertion {style-1 | style-2 | disable}
    set dhcp-option82-remote-id-insertion {style-1 | disable}
  next
end

```

Wireless network monitoring

You can monitor both your wireless clients and other wireless networks that are available in your coverage area.

Monitoring wireless clients

To view connected clients on a FortiWiFi unit

1. Go to **Monitor > Client Monitor**.

The following information is displayed:

SSID	The SSID that the client connected to.
FortiAP	The serial number of the FortiAP unit to which the client connected.
User	User name
IP	The IP address assigned to the wireless client.
Device	
Auth	The type of authentication used.
Channel	WiFi radio channel in use.
Bandwidth Tx/Rx	Client received and transmitted bandwidth, in Kbps.
Signal Strength / Noise	The signal-to-noise ratio in deciBels calculated from signal strength and noise level.
Signal Strength	
Association Time	How long the client has been connected to this access point.

Results can be filtered. Select the filter icon on the column you want to filter. Enter the values to include or select NOT if you want to exclude the specified values.

Monitoring rogue APs

The access point radio equipment can scan for other available access points, either as a dedicated monitor or in idle periods during AP operation.

Discovered access points are listed in **Monitor > Rogue AP Monitor**. You can then mark them as either Accepted or Rogue access points. This designation helps you to track access points. It does not affect anyone's ability to use these access points.

It is also possible to suppress rogue APs. See [Monitoring rogue APs on page 1145](#).

On-wire rogue AP detection technique

Other APs that are available in the same area as your own APs are not necessarily rogues. A neighboring AP that has no connection to your network might cause interference, but it is not a security threat. A rogue AP is an unauthorized AP connected to your wired network. This can enable unauthorized access. When rogue AP detection is enabled, the **On-wire** column in the **Rogue AP Monitor** list shows a green up-arrow on detected rogues.

Rogue AP monitoring of WiFi client traffic builds a table of WiFi clients and the Access Points that they are communicating through. The FortiGate unit also builds a table of MAC addresses that it sees on the LAN. The FortiGate unit's on-wire correlation engine constantly compares the MAC addresses seen on the LAN to the MAC addresses seen on the WiFi network.

There are two methods of Rogue AP on-wire detection operating simultaneously: Exact MAC address match and MAC adjacency.

Exact MAC address match

If the same MAC address is seen on the LAN and on the WiFi network, this means that the wireless client is connected to the LAN. If the AP that the client is using is not authorized in the FortiGate unit configuration, that AP is deemed an 'on-wire' rogue. This scheme works for non-NAT rogue APs.

MAC adjacency

If an access point is also a router, it applies NAT to WiFi packets. This can make rogue detection more difficult. However, an AP's WiFi interface MAC address is usually in the same range as its wired MAC address. So, the MAC adjacency rogue detection method matches LAN and WiFi network MAC addresses that are within a defined numerical distance of each other. By default, the MAC adjacency value is 7. If the AP for these matching MAC addresses is not authorized in the FortiGate unit configuration, that AP is deemed an 'on-wire' rogue.

Limitations

On-wire rogue detection has some limitations. There must be at least one WiFi client connected to the suspect AP and continuously sending traffic. If the suspect AP is a router, its WiFi MAC address must be very similar to its Ethernet port MAC address.

Logging

Information about detected rogue APs is logged and uploaded to your FortiAnalyzer unit, if you have one. By default, rogue APs generate an alert level log, unknown APs generate a warning level log. This log information can help you with PCI-DSS compliance requirements.

Rogue AP scanning as a background activity

Each WiFi radio can perform monitoring of radio channels in its operating band while acting as an AP. It does this by briefly switching from AP to monitoring mode. By default, a scan period starts every 300 seconds. Each second a different channel is monitored for 20ms until all channels have been checked.

During heavy AP traffic, it is possible for Spectrum Analysis background scanning to cause lost packets when the radio switches to monitoring. To reduce the probability of lost packets, you can set the CLI `ap-bgscan-idle` field to delay the switch to monitoring until the AP has been idle for a specified period. This means that heavy AP traffic may slow background scanning.

The following CLI example configures default background rogue scanning operation except that it sets `ap-bgscan-idle` to require 100ms of AP inactivity before scanning the next channel.

```
config wireless-controller wtp-profile
  edit ourprofile
    config radio-1
      set wids-profile ourwidsprofile
      set spectrum-analysis enable
    end
  end
config wireless-controller wids-profile
  edit ourwidsprofile
    set ap-scan enable
    set rogue-scan enable
    set ap-bgscan-period 300
    set ap-bgscan-intv 1
    set ap-bgscan-duration 20
    set ap-bgscan-idle 100
  end
```

Configuring rogue scanning

All APs using the same FortiAP Profile share the same rogue scanning settings, unless override is configured.

To enable rogue AP scanning with on-wire detection - web-based manager

1. Go to **WiFi & Switch Controller > WIDS Profiles**.
On some models, the menu is **WiFi & Switch Controller**.
2. Select an existing WIDS Profile and edit it, or select **Create New**.
3. Make sure that **Enable Rogue AP Detection** is selected.
4. Select **Enable On-Wire Rogue AP Detection**.
5. Optionally, enable **Auto Suppress Rogue APs in Foreground Scan**.
6. Select **OK**.

To enable the rogue AP scanning feature in a custom AP profile - CLI

```
config wireless-controller wids-profile
  edit FAP220B-default
    set ap-scan enable
    set rogue-scan enable
  end
```

Exempting an AP from rogue scanning

By default, if Rogue AP Detection is enabled, it is enabled on all managed FortiAP units. Optionally, you can exempt an AP from scanning. You should be careful about doing this if your organization must perform scanning to meet PCI-DSS requirements.

To exempt an AP from rogue scanning

1. Go to **WiFi & Switch Controller > WIDS Profiles**.
2. Create a new WIDS profile and disable **Rogue AP detection**.

3. Go to **WiFi & Switch Controller > FortiAP Profiles** and edit the profile you wish to exempt from rogue scanning.
4. Assign the WIDS profile created in step 2.

MAC adjacency

You can adjust the maximum WiFi to Ethernet MAC difference used when determining whether an suspect AP is a rogue.

To adjust MAC adjacency

For example, to change the adjacency to 8, enter








```
config wireless-controller global
  set rogue-scan-mac-adjacency 8
end
```

Using the Rogue AP Monitor

Go to **Monitor > Rogue AP Monitor** to view the list of other wireless access points that are receivable at your location.

Information Columns

Actual columns displayed depends on **Column Settings**.

State	 Rogue AP — Use this status for unauthorized APs that On-wire status indicates are attached to your wired networks.
	 Accepted AP — Use this status for APs that are an authorized part of your network or are neighboring APs that are not a security threat. To see accepted APs in the list, select Show Accepted .
	 Unclassified — This is the initial status of a discovered AP. You can change an AP back to unclassified if you have mistakenly marked it as Rogue or Accepted.
Online Status	 Active AP
	 Inactive AP
	 Active ad-hoc WiFi device
	 Inactive ad-hoc WiFi device
SSID	The wireless service set identifier (SSID) or network name for the wireless interface.
Security Type	The type of security currently being used.
Channel	The wireless radio channel that the access point uses.
MAC Address	The MAC address of the Wireless interface.
Vendor Info	The name of the vendor.
Signal Strength	The relative signal strength of the AP. Mouse over the symbol to view the signal-to-noise ratio.
Detected By	The name or serial number of the AP unit that detected the signal.
On-wire	A green up-arrow indicates a suspected rogue, based on the on-wire detection technique. A red down-arrow indicates AP is not a suspected rogue.
First Seen	How long ago this AP was first detected.

Last Seen	How long ago this AP was last detected.
Rate	Data rate in bps.

To change the Online Status of an AP, right-click it and select **Mark Accepted** or **Mark Rogue**.

Suppressing rogue APs

In addition to monitoring rogue APs, you can actively prevent your users from connecting to them. When suppression is activated against an AP, the FortiGate WiFi controller sends deauthentication messages to the rogue AP's clients, posing as the rogue AP, and also sends deauthentication messages to the rogue AP, posing as its clients. This is done using the monitoring radio.



Before enabling this feature, verify that operation of Rogue Suppression is compliant with the applicable laws and regulations of your region.

To enable rogue AP suppression, you must enable monitoring of rogue APs with the on-wire detection technique. See "[Monitoring rogue APs](#)". The monitoring radio must be in the Dedicated Monitor mode.

To activate AP suppression against a rogue AP

1. Go to **Monitor > Rogue AP Monitor**.
2. When you see an AP listed that is a rogue detected "on-wire", select it and then select **Mark > Mark Rogue**.
3. To suppress an AP that is marked as a rogue, select it and then select **Suppress AP**.

To deactivate AP suppression

1. Go to **Monitor > Rogue AP Monitor**.
2. Select the suppressed rogue AP and then select **Suppress AP > Unsuppress AP**.

Monitoring wireless network health

To view the wireless health dashboard, go to **Monitor > WiFi Health Monitor**.

The wireless health dashboard provides a comprehensive view of the health of your network's wireless infrastructure. The dashboard includes widgets to display:

- **AP Status**
Active, Down or missing, up for over 24 hours, rebooted in past 24 hours
- **Client Count Over Time**
Viewable for past hour, day, or 30 days
- **Top Client Count Per-AP**
Separate widgets for 2.4GHz and 5GHz bands
- **Top Wireless Interference**
Separate widgets for 2.4GHz and 5GHz bands, requires spectrum analysis to be enabled on the radios

- **Login Failures Information**
- **WiFi Channel Utilization**

Three views allowing users to view top 10-20 Most and Least utilized channels for each AP radio and a third histogram view showing counts for utilization

The list of active clients also shows MAC address entries (similar to the **WiFi Client Monitor** page), making client information easy to view when opening the **Active Client** widget.

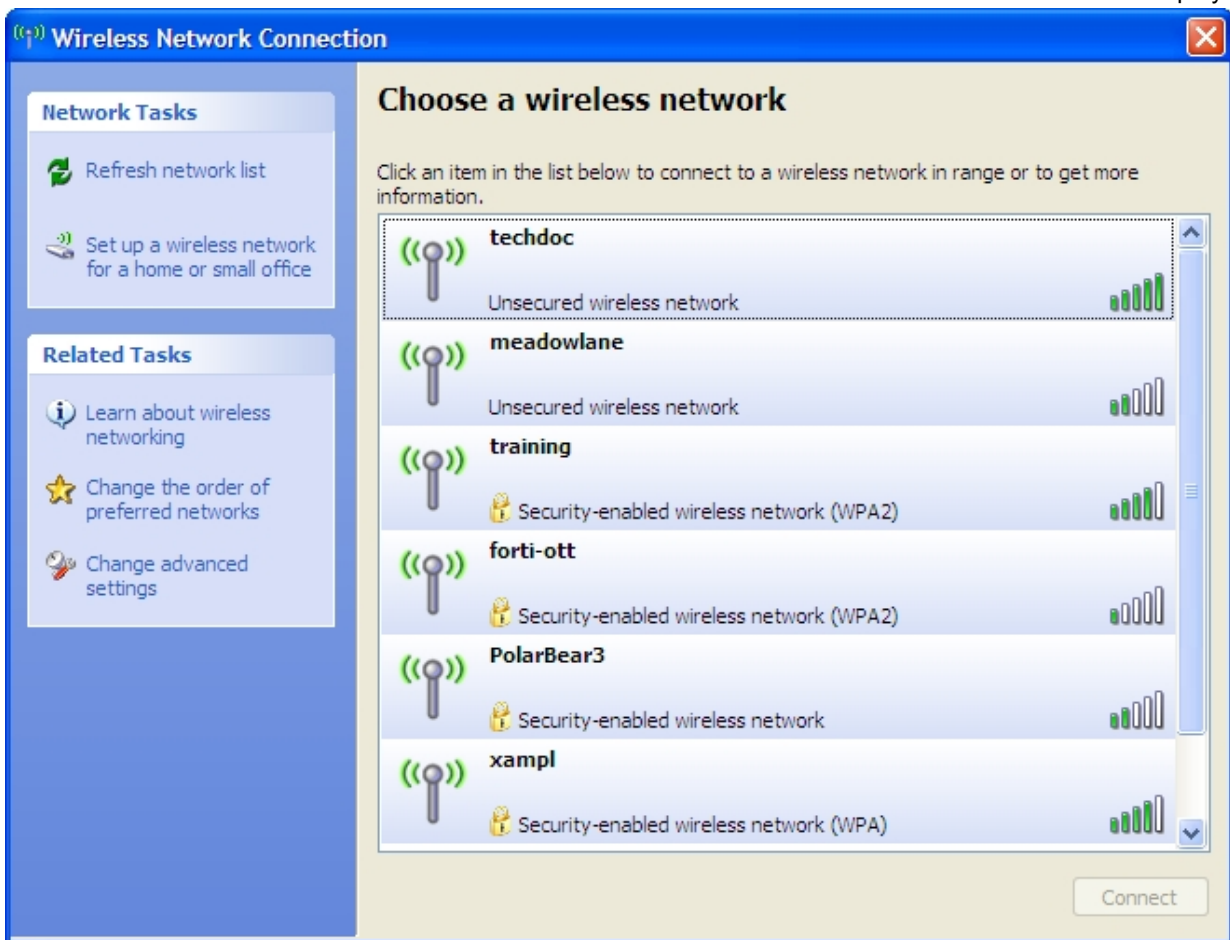
Configuring wireless network clients

This chapter shows how to configure typical wireless network clients to connect to a wireless network with WPA-Enterprise security.

Windows XP client

To configure the WPA-Enterprise network connection

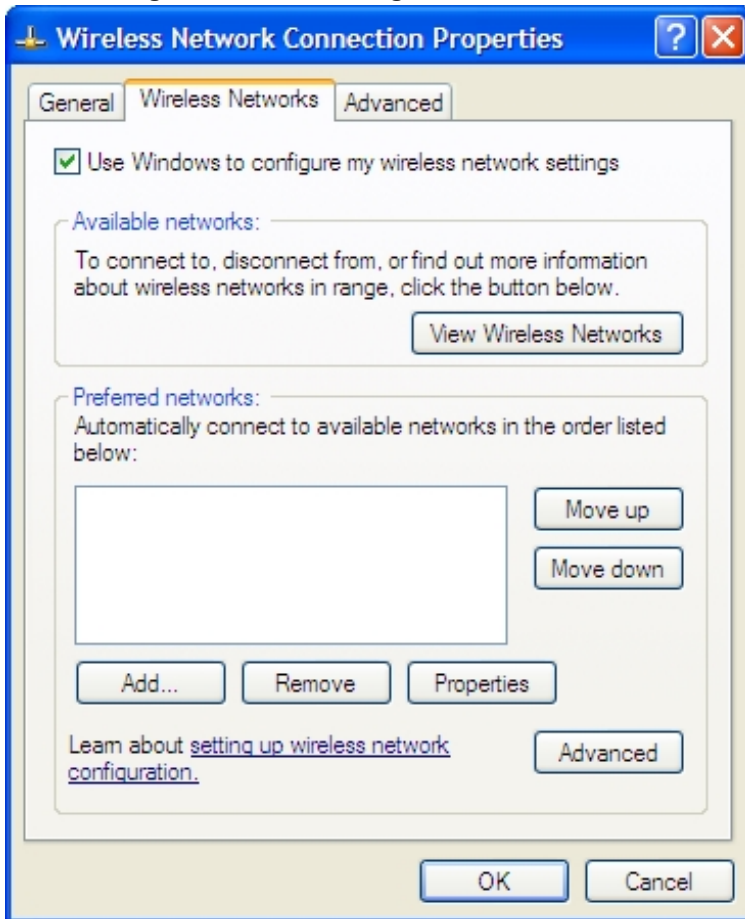
1. In the Windows Start menu, go to **Control Panel > Network Connections > Wireless Network Connection** or select the wireless network icon in the Notification area of the Taskbar. A list of available networks is displayed.



If you are already connected to another wireless network, the Connection Status window displays. Select **View Wireless Networks** on the **General** tab to view the list.

If the network broadcasts its SSID, it is listed. But do not try to connect until you have completed the configuration step below. Because the network doesn't use the Windows XP default security configuration, configure the client's network settings manually before trying to connect.

2. You can configure the WPA-Enterprise network to be accessible from the **View Wireless Networks** window even if it does not broadcast its SSID.
3. Select **Change Advanced Settings** and then select the **Wireless Networks** tab.



Any existing networks that you have already configured are listed in the **Preferred Networks** list.

4. Select **Add** and enter the following information:

Wireless network properties

Association Authentication Connection

Network name (SSID): xample

☐ Connect even if this network is not broadcasting

Wireless network key

This network requires a key for the following:

Network Authentication: WPA2

Data encryption: AES

Network key:

Confirm network key:

Key index (advanced): 1

☐ The key is provided for me automatically

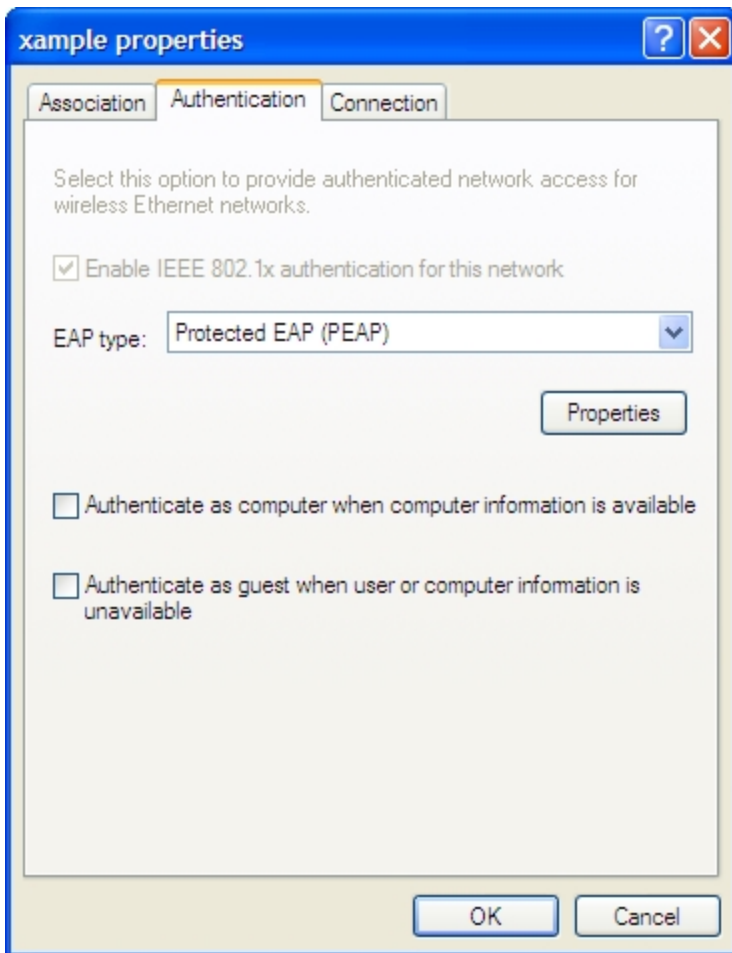
☐ This is a computer-to-computer (ad hoc) network; wireless access points are not used

OK Cancel

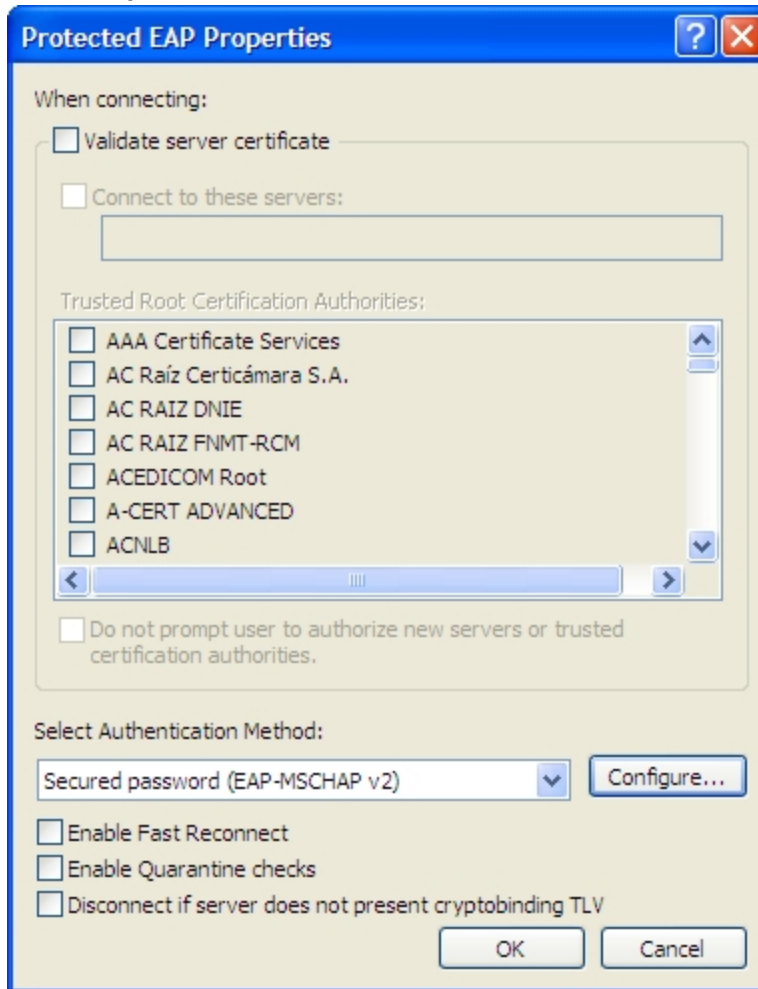
Network Name (SSID)	The SSID for your wireless network
Network Authentication	WPA2
Data Encryption	AES

5. If this wireless network does not broadcast its SSID, select **Connect even if this network is not broadcasting** so that the network will appear in the **View Wireless Networks** list.

6. Select the **Authentication** tab.



7. In **EAP Type**, select **Protected EAP (PEAP)**.
8. Make sure that the other two authentication options are not selected.

9. Select **Properties**.

10. Make sure that **Validate server certificate** is selected.
11. Select the server certificate **Entrust Root Certification Authority**.
12. In **Select Authentication Method**, select **Secured Password (EAP-MSCHAPv2)**.
13. Ensure that the remaining options are not selected.
14. Select **Configure**.



15. If your wireless network credentials are the same as your Windows logon credentials, select **Automatically use my Windows logon name and password**. Otherwise, make sure that this option is not selected.
16. Select **OK**. Repeat until you have closed all of the **Wireless Network Connection Properties** windows.

To connect to the WPA-Enterprise wireless network

1. Select the wireless network icon in the Notification area of the Taskbar.
2. In the **View Wireless Networks** list, select the network you just added and then select **Connect**. You might need to log off of your current wireless network and refresh the list.
3. When the following popup displays, click on it.




4. In the **Enter Credentials** window, enter your wireless network **User name**, **Password**, and **Logon domain** (if applicable). Then, select **OK**.



In future, Windows will automatically send your credentials when you log on to this network.

Windows 7 client

1. In the Windows Start menu, go to **Control Panel > Network and Internet > Network and Sharing Center > Manage Wireless Networks** or select the wireless network icon in the Notification area of the Taskbar. A list of available networks is displayed.



2. Do one of the following:
 - If the wireless network is listed (it broadcasts its SSID), select it from the list.
 - Select **Add > Manually create a network profile**.

3. Enter the following information and select **Next**.



Manually connect to a wireless network

Enter information for the wireless network you want to add

Network name:

Security type:

Encryption type:

Security Key: ☐ Hide characters

☒ Start this connection automatically

☒ Connect even if the network is not broadcasting

Warning: If you select this option, your computer's privacy might be at risk.

Network name	Enter the SSID of the wireless network. (Required only if you selected Add.)
Security type	WPA2-Enterprise
Encryption type	AES
Start this connection automatically	Select
Connect even if the network is not broadcasting.	Select

The Wireless Network icon will display a popup requesting that you click to enter credentials for the network. Click on the popup notification.

4. In the **Enter Credentials** window, enter your wireless network **User name**, **Password**, and **Logon domain** (if applicable). Then, select **OK**.
5. Select **Change connection settings**.
6. On the **Connection** tab, select **Connect automatically when this network is in range**.
7. On the **Security** tab, select the Microsoft PEAP authentication method and then select **Settings**.

8. Make sure that **Validate server certificate** is selected.
9. Select the server certificate **Entrust Root Certification Authority**.
10. In **Select Authentication Method**, select **Secured Password (EAP-MSCHAPv2)**.
11. Select **Configure**.




12. If your wireless network credentials are the same as your Windows logon credentials, select **Automatically use my Windows logon name and password**. Otherwise, make sure that this option is not selected.
13. Ensure that the remaining options are not selected.
14. Select **OK**. Repeat until you have closed all of the **Wireless Network Properties** windows.

Mac OS client

To configure network preferences

1. Right-click the **AirPort** icon in the toolbar and select **Open Network Preferences**.
2. Select **Advanced** and then select the **802.1X** tab.




3. If there are no Login Window Profiles in the left column, select the + button and then select **Add Login Window Profile**.
4. Select the Login Window Profile and then make sure that both TTLS and PEAP are selected in **Authentication**.

To configure the WPA-Enterprise network connection

1. Select the **AirPort** icon in the toolbar.
2. Do one of the following:
 - If the network is listed, select the network from the list.
 - Select **Connect to Other Network**.

One of the following windows opens, depending on your selection.




The network "xampl" requires a password.

User Name:

Password:

802.1X:

☒ Remember this network



Enter the name of the network.

Enter the name of the network you want to join, and then enter the password if necessary.

Network Name:

Security:

User Name:

Password:

802.1X:

☒ Remember this network

3. Enter the following information and select **OK** or **Join**:

Network name	Enter the SSID of your wireless network. (Other network only)
Wireless Security	WPA Enterprise
802.1X	Automatic
Username Password	Enter your logon credentials for the wireless network.
Remember this network	Select.

You are connected to the wireless network.



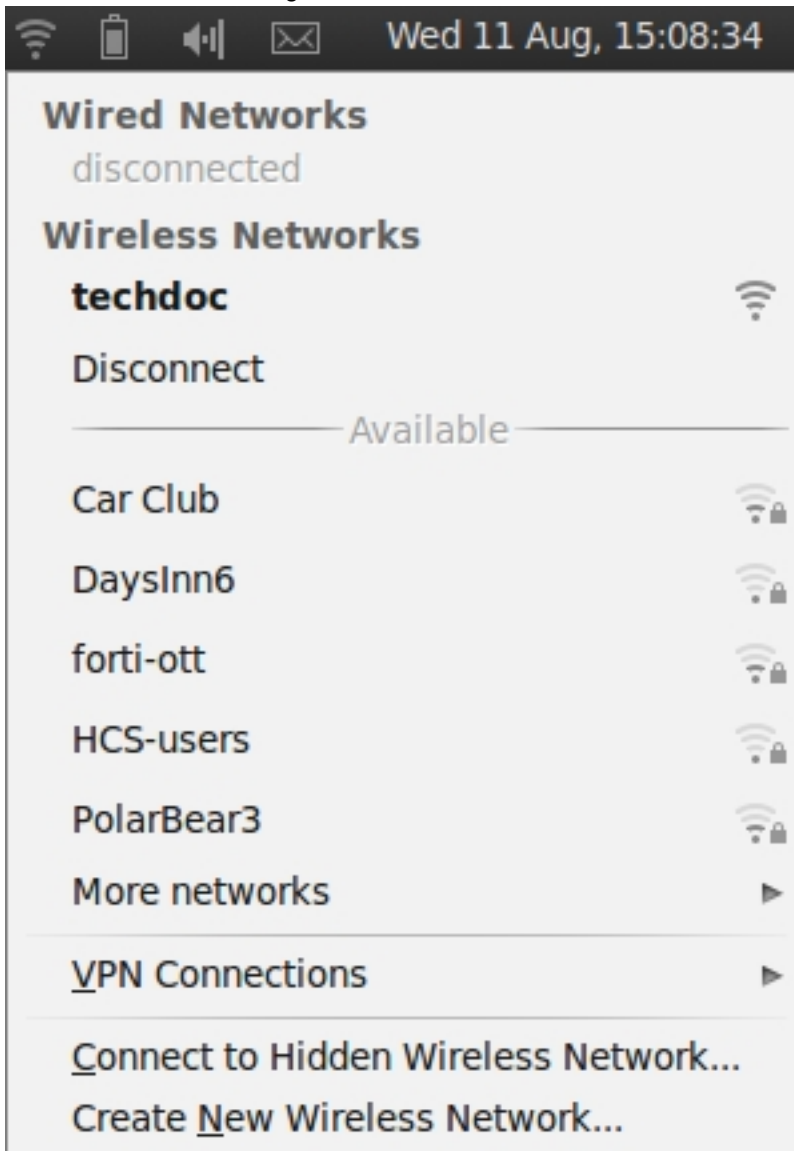
Mac OS supports only PEAP with MSCHAPv2 authentication and therefore can authenticate only to a RADIUS server, not an LDAP or TACACS+ server

Linux client

This example is based on the Ubuntu 10.04 Linux wireless client.

To connect to a WPA-Enterprise network

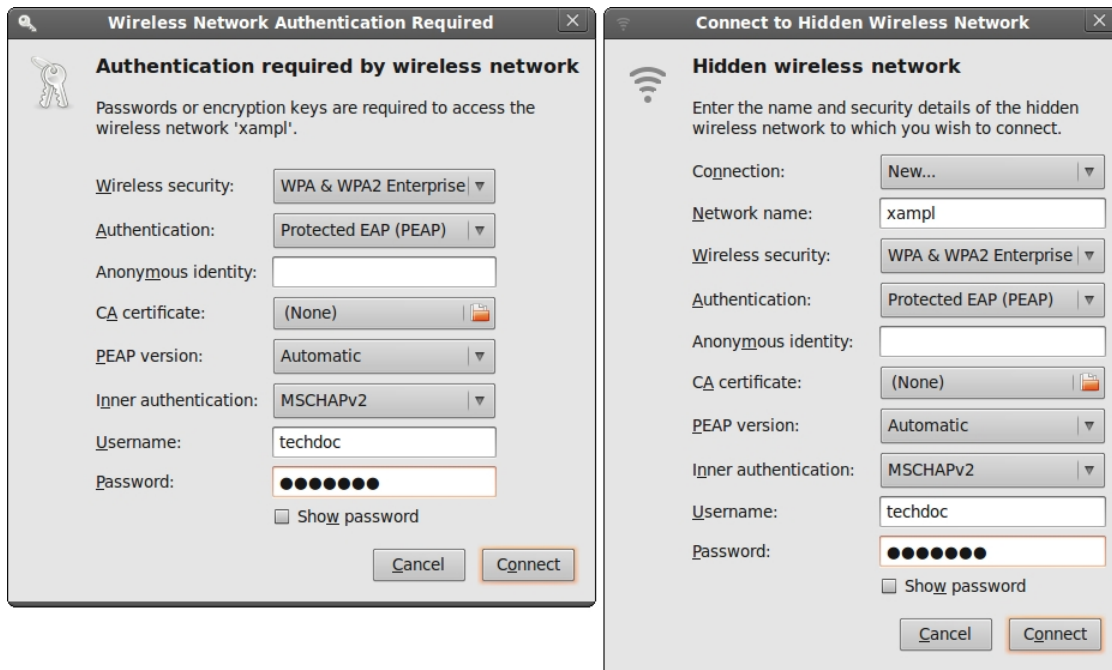
1. Select the Network Manager icon to view the Wireless Networks menu.



Wireless networks that broadcast their SSID are listed in the **Available** section of the menu. If the list is long, it is continued in the **More Networks** submenu.

2. Do one of the following:
 - Select the network from the list (also check **More Networks**).
 - Select **Connect to Hidden Wireless Network**.

One of the following windows opens, depending on your selection.



3. Enter the following information:

Connection	Leave as New . (Hidden network only)
Network name	Enter the SSID of your wireless network. (Hidden network only)
Wireless Security	WPA & WPA2 Enterprise
Authentication	Protected EAP (PEAP) for RADIUS-based authentication Tunneled TLS for TACACS+ or LDAP-based authentication
Anonymous identity	This is not required.
CA Certificate	If you want to validate the AP's certificate, select the Entrust Root Certification Authority root certificate. The default location for the certificate is /usr/share/ca-certificates/mozilla/.
PEAP version	Automatic (applies only to PEAP)
Inner authentication	MSCHAPv2 for RADIUS-based authentication PAP or CHAP for TACACS+ or LDAP-based authentication
Username Password	Enter your logon credentials for the wireless network.

4. If you did not select a CA Certificate above, you are asked to do so. Select Ignore.



5. Select **Connect**. You are connected to the wireless network.

To connect to a WPA-Enterprise network

1. Select the Network Manager icon to view the Wireless Networks menu.
2. Select the network from the list (also check **More Networks**).
If your network is not listed (but was configured), select **Connect to Hidden Wireless Network**, select your network from the Connection drop-down list, and then select **Connect**.

Troubleshooting

Using tools provided in your operating system, you can find the source of common wireless networking problems.

Checking that client received IP address and DNS server information

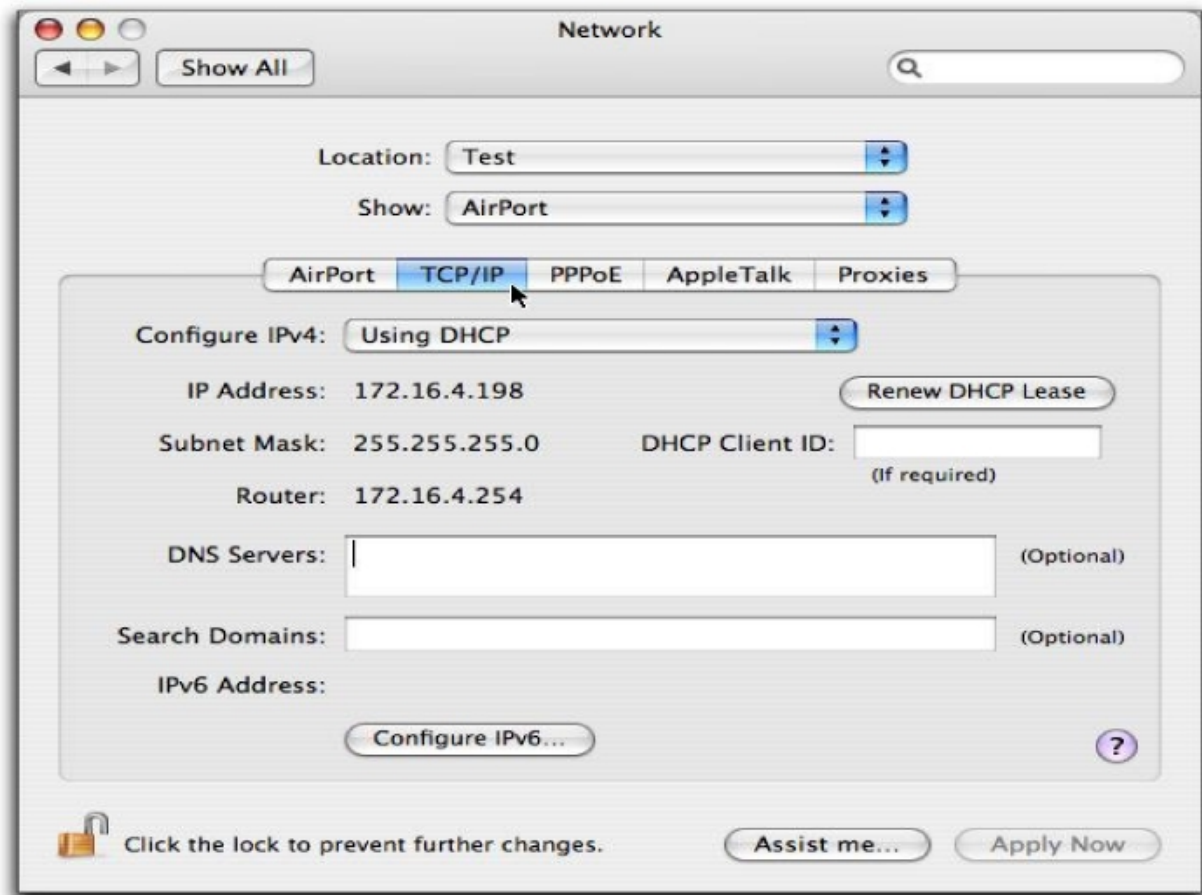
Windows XP

1. Double-click the network icon in the taskbar to display the **Wireless Network Connection Status** window. Check that the correct network is listed in the **Connection** section.
2. Select the **Support** tab.
Check that the **Address Type** is **Assigned by DHCP**. Check that the **IP Address**, **Subnet Mask**, and **Default Gateway** values are valid.
3. Select **Details** to view the DNS server addresses.
The listed address should be the DNS serves that were assigned to the WAP. Usually a wireless network that provides access to the private LAN is assigned the same DNS servers as the wired private LAN. A wireless network that provides guest or customer users access to the Internet is usually assigned public DNS servers.
4. If any of the addresses are missing, select **Repair**.
If the repair procedure doesn't correct the problem, check your network settings.

Mac OS

1. From the Apple menu, open **System Preferences > Network**.
2. Select **AirPort** and then select **Configure**.

- On the **Network** page, select the **TCP/IP** tab.



- If there is no IP address or the IP address starts with 169, select **Renew DHCP Lease**.
- To check DNS server addresses, open a terminal window and enter the following command:

```
cat /etc/resolv.conf
```

Check the listed nameserver addresses. A network for employees should use the wired private LAN DNS server. A network for guests should specify a public DNS server.

Linux

This example is based on the Ubuntu 10.04 Linux wireless client.

1. Right-click the Network Manager icon and select **Connection Information**.



2. Check the IP address, and DNS settings. If they are incorrect, check your network settings.

Wireless network examples

This chapter provides an example wireless network configuration.

Basic wireless network

This example uses automatic configuration to set up a basic wireless network.

To configure this wireless network, you must:

- Configure authentication for wireless users
- Configure the SSID (WiFi network interface)
- Add the SSID to the FortiAP Profile
- Configure the firewall policy
- Configure and connect FortiAP units

Configuring authentication for wireless users

You need to configure user accounts and add the users to a user group. This example shows only one account, but multiple accounts can be added as user group members.

To configure a WiFi user - web-based manager

1. Go to **User & Device > User Definition** and select **Create New**.
2. Select **Local User** and then click **Next**.
3. Enter a **User Name** and **Password** and then click **Next**.
4. Click **Next**.
5. Make sure that **Enable** is selected and then click **Create**.

To configure the WiFi user group - web-based manager

1. Go to **User & Device > User Groups** and select **Create New**.
2. Enter the following information and then select **OK**:

Name	wlan_users
Type	Firewall
Members	Add users.

To configure a WiFi user and the WiFi user group - CLI

```
config user user
  edit "user01"
    set type password
    set passwd "asdf12ghjk"
  end
config user group
```

```
edit "wlan_users"
  set member "user01"
end
```

Configuring the SSID

First, establish the SSID (network interface) for the network. This is independent of the number of physical access points that will be deployed. The network assigns IP addresses using DHCP.

To configure the SSID - web-based manager

1. Go to **WiFi & Switch Controller > SSID** and select **Create New > SSID**.
2. Enter the following information and select **OK**:

Interface Name	example_wifi_if
Traffic Mode	Tunnel to Wireless Controller
IP/Network Mask	10.10.110.1/24
Administrative Access	Ping (to assist with testing)
DHCP Server	Enable
Address Range	10.10.110.2 - 10.10.110.199
Netmask	255.255.255.0
Default Gateway	Same As Interface IP
DNS Server	Same as System DNS
SSID	example_wifi
Security Mode	WPA2 Enterprise
Authentication	Local, select wlan_users user group.
Leave other settings at their default values.	

To configure the SSID - CLI

```
config wireless-controller vap
  edit example_wifi_if
    set ssid "example_wifi"
    set broadcast-ssid enable
    set security wpa-enterprise
    set auth usergroup
    set usergroup wlan_users
    set schedule always
  end
config system interface
  edit example_wifi_if
    set ip 10.10.110.1 255.255.255.0
  end
config system dhcp server
```



```

edit 0
    set default-gateway 10.10.110.1
    set dns-service default
    set interface "example_wifi_if"
    config ip-range
        edit 1
            set end-ip 10.10.110.199
            set start-ip 10.10.110.2
        end
    set netmask 255.255.255.0
end

```

Adding the SSID to the FortiAP Profile

The radio portion of the FortiAP configuration is contained in the FortiAP Profile. By default, there is a profile for each platform (FortiAP model). You can create additional profiles if needed. The SSID needs to be specified in the profile.

To add the SSID to the FortiAP Profile - web-based manager

1. Go to **WiFi & Switch Controller > FortiAP Profiles** and edit the profile for your model of FortiAP unit.
2. In **Radio 1** and **Radio 2**, add example_wifi in **SSID**.
3. Select **OK**.

Configuring security policies

A security policy is needed to enable WiFi users to access the Internet on port1. First you create firewall address for the WiFi network, then you create the example_wifi to port1 policy.

To create a firewall address for WiFi users - web-based manager

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New > Address**, enter the following information and select **OK**.

Name	wlan_user_net
Type	IP/Netmask
Subnet / IP Range	10.10.110.0/24
Interface	example_wifi_if
Show in Address List	Enabled

To create a firewall address for WiFi users - CLI

```

config firewall address
    edit "wlan_user_net"
        set associated-interface "example_wifi_if"
        set subnet 10.10.110.0 255.255.255.0
    end

```

To create a security policy for WiFi users - web-based manager

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information and select **OK**:

Incoming Interface	example_wifi_if
Source Address	wlan_user_net
Outgoing Interface	port1
Destination Address	All
Schedule	always
Service	ALL
Action	ACCEPT
NAT	ON. Select Use Destination Interface Address (default).
Leave other settings at their default values.	

To create a firewall policy for WiFi users - CLI

```
config firewall policy
  edit 0
    set srcintf "example_wifi"
    set dstintf "port1"
    set srcaddr "wlan_user_net"
    set dstaddr "all"
    set schedule always
    set service ALL
    set action accept
    set nat enable
  end
```

Connecting the FortiAP units

You need to connect each FortiAP unit to the FortiGate unit, wait for it to be recognized, and then assign it to the AP Profile. But first, you must configure the interface to which the FortiAP units connect and the DHCP server that assigns their IP addresses.

In this example, the FortiAP units connect to port 3 and are controlled through IP addresses on the 192.168.8.0/24 network.

To configure the interface for the AP unit - web-based manager

1. Go to **Network > Interfaces** and edit the port3 interface.
2. Set the **Addressing mode** to **Dedicated to Extension Device** and set the **IP/Network Mask** to 192.168.8.1/255.255.255.0.
3. Select **OK**.

This procedure automatically configures a DHCP server for the AP units.

To configure the interface for the AP unit - CLI

```
config system interface
edit port3
set mode static
set ip 192.168.8.1 255.255.255.0
end
```

To configure the DHCP server for AP units - CLI

```
config system dhcp server
edit 0
set interface port3
config exclude-range
edit 1
set end-ip 192.168.8.1
set start-ip 192.168.8.1
end
config ip-range
edit 1
set end-ip 192.168.8.254
set start-ip 192.168.8.2
end
set netmask 255.255.255.0
set vci-match enable
set vci-string "FortiAP"
end
```

To connect a FortiAP unit - web-based manager

1. Go to **WiFi & Switch Controller > Managed FortiAPs**.
2. Connect the FortiAP unit to port 3.
3. Periodically select **Refresh** while waiting for the FortiAP unit to be listed.
Recognition of the FortiAP unit can take up to two minutes.
If FortiAP units are connected but cannot be recognized, try disabling VCI-Match in the DHCP server settings.
4. When the FortiAP unit is listed, select the entry to edit it.
The **Edit Managed Access Point** window opens.
5. In **State**, select **Authorize**.
6. In **FortiAP Profile**, select the default profile for the FortiAP model.
7. Select **OK**.
8. Repeat Steps 2 through 8 for each FortiAP unit.

To connect a FortiAP unit - CLI

1. Connect the FortiAP unit to port 3.
2. Enter

```
config wireless-controller wtp
```

3. Wait 30 seconds, then enter `get`.

Retry the `get` command every 15 seconds or so until the unit is listed, like this:

```
== [ FAP22B3U10600118 ]
wtp-id: FAP22B3U10600118
```

4. Edit the discovered FortiAP unit like this:

```
edit FAP22B3U10600118
    set admin enable
end
```

5. Repeat Steps 2 through 4 for each FortiAP unit.

A more complex example

This example creates multiple networks and uses custom AP profiles.

Scenario

In this example, Example Co. provides two wireless networks, one for its employees and the other for customers or other guests of its business. Guest users have access only to the Internet, not to the company's private network. The equipment for these WiFi networks consists of FortiAP-220B units controlled by a FortiGate unit.

The employee network operates in 802.11n mode on both the 2.4GHz and 5GHz bands. Client IP addresses are in the 10.10.120.0/24 subnet, with 10.10.120.1 the IP address of the WAP. The guest network also operates in 802.11n mode, but only on the 2.4GHz band. Client IP addresses are on the 10.10.115.0/24 subnet, with 10.10.115.1 the IP address of the WAP.

On FortiAP-220B units, the 802.11n mode also supports 802.11g and 802.11b clients on the 2.4GHz band and 802.11a clients on the 5GHz band.

The guest network WAP broadcasts its SSID, the employee network WAP does not.

The employees network uses WPA-Enterprise authentication through a FortiGate user group. The guest network features a captive portal. When a guest first tries to connect to the Internet, a login page requests logon credentials. Guests use numbered guest accounts authenticated by RADIUS. The captive portal for the guests includes a disclaimer page.

In this example, the FortiAP units connect to port 3 and are assigned addresses on the 192.168.8.0/24 subnet.

Configuration

To configure these wireless networks, you must:

- Configure authentication for wireless users
- Configure the SSIDs (network interfaces)
- Configure the AP profile
- Configure the WiFi LAN interface and a DHCP server
- Configure firewall policies

Configuring authentication for employee wireless users

Employees have user accounts on the FortiGate unit. This example shows creation of one user account, but you can create multiple accounts and add them as members to the user group.

To configure a WiFi user - web-based manager

1. Go to **User & Device > User Definition** and select **Create New**.
2. Select **Local User** and then click **Next**.

3. Enter a **User Name** and **Password** and then click **Next**.
4. Click **Next**.
5. Make sure that **Enable** is selected and then click **Create**.

To configure the user group for employee access - web-based manager

1. Go to **User & Device > User Groups** and select **Create New**.
2. Enter the following information and then select **OK**:

Name	employee-group
Type	Firewall
Members	Add users.

To configure a WiFi user and the user group for employee access - CLI

```
config user user
  edit "user01"
    set type password
    set passwd "asdf12ghjk"
  end
config user group
  edit "employee-group"
    set member "user01"
  end
```

The user authentication setup will be complete when you select the employee-group in the SSID configuration.

Configuring authentication for guest wireless users

Guests are assigned temporary user accounts created on a RADIUS server. The RADIUS server stores each user's group name in the Fortinet-Group-Name attribute. Wireless users are in the group named "wireless".

The FortiGate unit must be configured to access the RADIUS server.

To configure the FortiGate unit to access the guest RADIUS server - web-based manager

1. Go to **User & Device > RADIUS Servers** and select **Create New**.
2. Enter the following information and select OK:

Name	guestRADIUS
Primary Server IP/Name	10.11.102.100
Primary Server Secret	grikfwpdfg
Secondary Server IP/Name	Optional
Secondary Server Secret	Optional
Authentication Scheme	Use default, unless server requires otherwise.
Leave other settings at their default values.	

To configure the FortiGate unit to access the guest RADIUS server - CLI

```
config user radius
  edit guestRADIUS
    set auth-type auto
    set server 10.11.102.100
    set secret grikfwpfdg
  end
```

To configure the user group for guest access - web-based manager

1. Go to **User & Device > User Groups** and select **Create New**.
2. Enter the following information and then select **OK**:

Name	guest-group
Type	Firewall
Members	Leave empty.

3. Select **Create new**.
4. Enter:

Remote Server	Select guestRADIUS .
Groups	Select wireless

5. Select **OK**.

To configure the user group for guest access - CLI

```
config user group
  edit "guest-group"
    set member "guestRADIUS"
    config match
      edit 0
        set server-name "guestRADIUS"
        set group-name "wireless"
      end
    end
  end
```

The user authentication setup will be complete when you select the guest-group user group in the SSID configuration.

Configuring the SSIDs

First, establish the SSIDs (network interfaces) for the employee and guest networks. This is independent of the number of physical access points that will be deployed. Both networks assign IP addresses using DHCP.

To configure the employee SSID - web-based manager

1. Go to **WiFi & Switch Controller > SSID** and select **Create New > SSID**.
2. Enter the following information and select **OK**:

Interface Name	example_inc
Traffic Mode	Tunnel to Wireless Controller
IP/Netmask	10.10.120.1/24
Administrative Access	Ping (to assist with testing)
Enable DHCP	Enable
Address Range	10.10.120.2 - 10.10.120.199
Netmask	255.255.255.0
Default Gateway	Same As Interface IP
DNS Server	Same as System DNS
SSID	example_inc
Security Mode	WPA/WPA2-Enterprise
Authentication	Select Local , then select employee-group .
Leave other settings at their default values.	

To configure the employee SSID - CLI

```

config wireless-controller vap
  edit example_inc
    set ssid "example_inc"
    set security wpa-enterprise
    set auth usergroup
    set usergroup employee-group
    set schedule always
  end
config system interface
  edit example_inc
    set ip 10.10.120.1 255.255.255.0
  end
config system dhcp server
  edit 0
    set default-gateway 10.10.120.1
    set dns-service default
    set interface example_inc
    config ip-range
      edit 1
        set end-ip 10.10.120.199
        set start-ip 10.10.120.2
      end
    set lease-time 7200
    set netmask 255.255.255.0
  end
end

```

To configure the example_guest SSID - web-based manager

1. Go to **WiFi & Switch Controller > SSID** and select **Create New**.
2. Enter the following information and select **OK**:

Name	example_guest
IP/Netmask	10.10.115.1/24
Administrative Access	Ping (to assist with testing)
Enable DHCP	Enable
Address Range	10.10.115.2 - 10.10.115.50
Netmask	255.255.255.0
Default Gateway	Same as Interface IP
DNS Server	Same as System DNS
SSID	example_guest
Security Mode	Captive Portal
Portal Type	Authentication
Authentication Portal	Local
User Groups	Select guest-group
Leave other settings at their default values.	

To configure the example_guest SSID - CLI

```
config wireless-controller vap
  edit example_guest
    set ssid "example_guest"
    set security captive-portal
    set selected-usergroups guest-group
    set schedule always
  end
config system interface
  edit example_guest
    set ip 10.10.115.1 255.255.255.0
  end
config system dhcp server
  edit 0
    set default-gateway 10.10.115.1
    set dns-service default
    set interface "example_guest"
    config ip-range
      edit 1
        set end-ip 10.10.115.50
        set start-ip 10.10.115.2
      end
    set lease-time 7200
```



```

set netmask 255.255.255.0
end

```

Configuring the FortiAP profile

The FortiAP Profile defines the radio settings for the networks. The profile provides access to both Radio 1 (2.4GHz) and Radio 2 (5GHz) for the employee virtual AP, but provides access only to Radio 1 for the guest virtual AP.

To configure the FortiAP Profile - web-based manager

1. Go to **WiFi & Switch Controller > FortiAP Profiles** and select **Create New**.
2. Enter the following information and select **OK**:

Name	example_AP
Platform	FAP220B
Radio 1	
Mode	Access Point
Band	802.11n
Channel	Select 1, 6, and 11.
Tx Power	100%
SSID	Select SSIDs and select example_inc and example_guest .
Radio 2	
Mode	Access Point
Band	802.11n_5G
Channel	Select all.
Tx Power	100%
SSID	Select SSIDs and select example_inc .

To configure the AP Profile - CLI

```

config wireless-controller wtp-profile
  edit "example_AP"
    config platform
      set type 220B
    end
    config radio-1
      set ap-bgscan enable
      set band 802.11n
      set channel "1" "6" "11"
      set vaps "example_inc" "example_guest"
    end
    config radio-2

```

```
set ap-bgscan enable
set band 802.11n-5G
set channel "36" "40" "44" "48" "149" "153" "157" "161" "165"
set vaps "example_inc"
end
```

Configuring firewall policies

Identity-based firewall policies are needed to enable the WLAN users to access the Internet on Port1. First you create firewall addresses for employee and guest users, then you create the firewall policies.

To create firewall addresses for employee and guest WiFi users

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**, enter the following information and select **OK**.

Address Name	employee-wifi-net
Type	Subnet / IP Range
Subnet / IP Range	10.10.120.0/24
Interface	example_inc

3. Select **Create New**, enter the following information and select **OK**.

Address Name	guest-wifi-net
Type	Subnet / IP Range
Subnet / IP Range	10.10.115.0/24
Interface	example_guest

To create firewall policies for employee WiFi users - web-based manager

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information and select **OK**:

Incoming Interface	example_inc
Source Address	employee-wifi-net
Outgoing Interface	port1
Destination Address	all
Schedule	always
Service	ALL
Action	ACCEPT
NAT	Enable NAT

3. Optionally, select security profile for wireless users.
4. Select **OK**.
5. Repeat steps 1 through 4 but select Internal as the Destination Interface/Zone to provides access to the ExampleCo private network.

To create firewall policies for employee WiFi users - CLI

```
config firewall policy
  edit 0
    set srcintf "employee_inc"
    set dstintf "port1"
    set srcaddr "employee-wifi-net"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY"
    set nat enable
    set schedule "always"
    set service "ANY"
  next
  edit 0
    set srcintf "employee_inc"
    set dstintf "internal"
    set srcaddr "employee-wifi-net"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY"
    set nat enable
    set schedule "always"
    set service "ANY"
  end
```

To create a firewall policy for guest WiFi users - web-based manager

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information and select **OK**:

Incoming Interface	example_guest
Source Address	guest-wifi-net
Outgoing Interface	port1
Destination Address	all
Schedule	always
Service	ALL
Action	ACCEPT
NAT	Enable NAT

3. Optionally, select **UTM** and set up UTM features for wireless users.
4. Select **OK**.

To create a firewall policy for guest WiFi users - CLI

```
config firewall policy
  edit 0
    set srcintf "example_guest"
    set dstintf "port1"
    set srcaddr "guest-wifi-net"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY"
    set nat enable
  end
```

Connecting the FortiAP units

You need to connect each FortiAP-220A unit to the FortiGate unit, wait for it to be recognized, and then assign it to the AP Profile. But first, you must configure the interface to which the FortiAP units connect and the DHCP server that assigns their IP addresses.

In this example, the FortiAP units connect to port 3 and are controlled through IP addresses on the 192.168.8.0/24 network.

To configure the interface for the AP unit - web-based manager

1. Go to **Network > Interfaces** and edit the port3 interface.
2. Set the **Addressing mode** to **Dedicated to Extension Device** and set the **IP/Netmask** to 192.168.8.1/255.255.255.0.
This step automatically configures a DHCP server for the AP units.
3. Select **OK**.

To configure the interface for the AP unit - CLI

```
config system interface
  edit port3
    set mode static
    set ip 192.168.8.1 255.255.255.0
  end
```

To configure the DHCP server for AP units - CLI

```
config system dhcp server
  edit 0
    set interface port3
    config ip-range
      edit 1
        set end-ip 192.168.8.9
        set start-ip 192.168.8.2
      end
    set netmask 255.255.255.0
    set vci-match enable
    set vci-string "FortiAP"
```

end

To connect a FortiAP-220A unit - web-based manager

1. Go to **WiFi & Switch Controller > Managed FortiAPs**.
2. Connect the FortiAP unit to port 3.
3. Periodically select **Refresh** while waiting for the FortiAP unit to be listed.
Recognition of the FortiAP unit can take up to two minutes.
If there is persistent difficulty recognizing FortiAP units, try disabling VCI-Match in the DHCP server settings.
4. When the FortiAP unit is listed, select the entry to edit it.
The **Edit Managed Access Point** window opens.
5. In **State**, select **Authorize**.
6. In the **AP Profile**, select **[Change]** and then select the **example_AP** profile.
7. Select **OK**.
8. Repeat Steps 2 through 8 for each FortiAP unit.

To connect a FortiAP-220A unit - CLI

1. Connect the FortiAP unit to port 3.
2. Enter:

```
config wireless-controller wtp
```
3. Wait 30 seconds, then enter `get`.
Retry the `get` command every 15 seconds or so until the unit is listed, like this:

```
== [ FAP22A3U10600118 ]
wtp-id: FAP22A3U10600118
```
4. Edit the discovered FortiAP unit like this:

```
edit FAP22A3U10600118
    set admin enable
    set wtp-profile example_AP
end
```
5. Repeat Steps 2 through 4 for each FortiAP unit.

Managing a FortiAP with FortiCloud

This chapter provides a few FortiCloud-managed FortiAP configuration examples.

You can register for a free FortiCloud account at www.forticloud.com.

For a video tutorial of how to configure and manage a FortiAP-S device from FortiCloud, follow the link below:

- [How to configure and Manage FortiAP-S from FortiCloud](#)

FortiCloud-managed FortiAP WiFi

In this example, you use FortiCloud to configure a single FortiAP-221C, creating a working WiFi network without a FortiGate.

FortiCloud remote management is supported on FortiAP models 221C and 320C.

For this configuration, the FortiAP-221C unit is running version 5.2 firmware. You will create a simple network that uses WPA-Personal authentication.

You can register for a free FortiCloud account at www.forticloud.com.

To create the WiFi network without a FortiGate unit, you must:

- Add your FortiAP to FortiCloud
- Configure the SSID
- Configure the AP platform profile
- Deploy the AP with the profile

Adding your FortiAP to FortiCloud

You need to add the FortiAP unit to your FortiCloud account. This is done through a unique key that can be found under the FortiAP unit.

To add a FortiAP to FortiCloud

1. Connect the FortiAP Ethernet interface to a network that provides access to the Internet.
2. Open a web browser and navigate to the FortiCloud main page and select **+ AP Network**.
3. Enter an **AP Network Name** and **AP Password**. This password is used to locally log in to the AP as the administrator. It will be set to all APs in this AP network.
4. Set the correct **Time Zone** and select **Submit**.

Configuring the SSID

You must establish the SSID (network interface) for the WiFi network.

To configure the SSID

1. Select the FortiAP you just created from the home page. You will then be prompted to add an SSID for the AP Network.
In the interface, this is under **Configure > SSIDs**.

2. In **Access Control**, enter the name of your SSID, set **Authentication** to **WPA2-Personal**, enter the **Pre-shared Key**, and select **Next**.
3. In **Security**, enable security features as required (select from **AntiVirus**, **Intrusion Prevention**, **Block Botnet**, **Web Access**, and **Application Control**) and select **Next**.
4. In **Availability**, make sure to leave **5 GHz** enabled, configure a schedule as required, and select **Next**.
5. Review your SSID in **Preview**, then select **Apply**.

Configuring the AP platform profile

The radio portion of the FortiAP configuration is contained in the FortiAP platform profile. By default, there is a profile for each platform (FortiAP model). The SSID needs to be specified in the profile.

To configure the AP profile

1. Go to **Configure > AP Profile** and edit the AP Profile for your FortiAP model (mouse-over the AP Profile to reveal the **Edit** button).
2. Enable the SSID configured earlier for both **Radio 1** and **Radio 2**, for 5GHz coverage.

Deploying the AP with the platform profile

With the SSID and platform profile configured, you must deploy the AP by entering the FortiCloud key for the FortiAP.

To deploy the AP

1. Go to **Configure > Deploy APs**. Here you will be prompted to enter the FortiCloud key, which can be found on the same label as the FortiAP unit's serial number, and select **Submit**.



If you have a FortiAP model that does *not* include a FortiCloud key, you can still add the device to the network. To learn how, see the [FortiCloud-managed FortiAP WiFi without a key](#) configuration.

2. In **Set Platform Profiles**, select the platform profile you created earlier and select **Next**.
3. Follow the rest of the deployment wizard. Select **Submit** when completed.

You will now be able to connect to the wireless network and browse the Internet. On the FortiCloud website, go to **Monitor > Report** where you can view monitoring information such as **Traffic by Period**, **Client Count by Period**, and more.

FortiCloud-managed FortiAP WiFi without a key

You can manage your FortiAP-based wireless network with FortiCloud even if your FortiAP has no FortiCloud key.

For this example, you will need to have already pre-configured your FortiAP unit with your FortiCloud account credentials. For more information on how to do this, or if your FortiAP has a FortiCloud key (on the serial number label), see the [FortiCloud-managed FortiAP WiFi](#) configuration.

You can register for a free FortiCloud account at www.forticloud.com.

To create the WiFi network without a FortiCloud key, you must:

- Configure the FortiAP unit
- Add the FortiAP unit to your FortiCloud account
- Configure the FortiAP

Configuring the FortiAP unit

You need to connect and configure the FortiAP unit through the web-based manager of the FortiGate.

To configure the FortiAP unit - web-based manager

1. Connect your computer to the FortiAP Ethernet port. The FortiAP's default IP address is 192.168.1.2. The computer should have an address on the same subnet, 192.168.1.3 for example.
2. Using a browser, log in to the FortiAP as *admin*. Leave the password field empty.
3. In **WTP-Configuration**, select **FortiCloud** and enter your FortiCloud credentials. Select **Apply**.
The FortiAP is now ready to connect to FortiCloud via the Internet.

Adding the FortiAP unit to your FortiCloud account

The FortiAP must be added to the FortiCloud account that has a WiFi network already configured for it.

For an example of creating a WiFi network on FortiCloud, see [FortiCloud-managed FortiAP WiFi on page 1183](#).

To add the FortiAP to FortiCloud

1. Connect the FortiAP Ethernet cable to a network that connects to the Internet.
Restore your computer to its normal network configuration and log on to FortiCloud.
2. From the **Home** screen, go to **Inventory > AP Inventory**. Your FortiAP should be listed.
3. Then go back to the Home screen, select your AP network, and go to **Deploy APs**.
4. Select your listed FortiAP and select **Next**.
5. Make sure your platform profile is selected from the dropdown menu, and select **Next**.
6. In **Preview**, select **Deploy**.
The device will now appear listed under **Access Points**.

You will now be able to connect to the wireless network and browse the Internet. On the FortiCloud website, go to **Monitor > Report** where you can view monitoring information such as **Traffic by Period**, **Client Count by Period**, and more.

Using a FortiWiFi unit as a client

A FortiWiFi operates by default as a wireless access point. But a FortiWiFi can also operate as a wireless client, connecting the FortiGate to another wireless network.

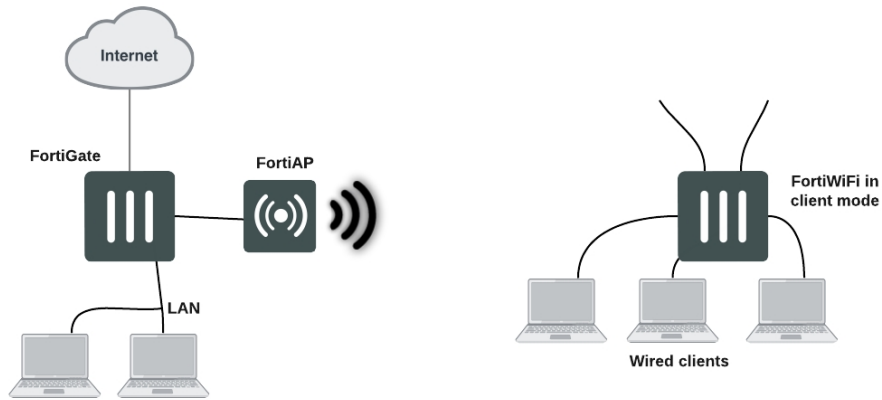
Use of client mode

In client mode, the FortiWiFi unit connects to a remote WiFi access point to access other networks or the Internet. This is most useful when the FortiWiFi unit is in a location that does not have a wired infrastructure.

For example, in a warehouse where shipping and receiving are on opposite sides of the building, running cables might not be an option due to the warehouse environment. The FortiWiFi unit can support wired users using its Ethernet ports and can connect to another access point wirelessly as a client. This connects the wired users to the network using the 802.11 WiFi standard as a backbone.

Note that in client mode the FortiWiFi unit cannot operate as an AP. WiFi clients cannot see or connect to the FortiWiFi unit in Client mode.

FortiWiFi unit in Client mode



Configuring client mode

To set up the FortiAP unit as a WiFi client, you must use the CLI. Before you do this, be sure to remove any AP WiFi configurations such as SSIDs, DHCP servers, policies, and so on.

To configure wireless client mode

1. Change the WiFi mode to client.

In the CLI, enter the following commands:

```
config system global
  set wireless-mode client
end
```

Respond "y" when asked if you want to continue. The FortiWiFi unit will reboot.

2. Configure the WiFi interface settings.

For example, to configure the client for WPA-Personal authentication on the **our_wifi** SSID with passphrase **justforus**, enter the following in the CLI:

```
config system interface
  edit wifi
    set mode dhcp
    config wifi-networks
      edit 0
        set wifi-ssid our_wifi
        set wifi-security wpa-personal
        set wifi-passphrase "justforus"
      end
    end
end
```

The WiFi interface client_wifi will receive an IP address using DHCP.

3. Configure a wifi to port1 policy.

You can use either CLI or web-based manager to do this. The important settings are:

Incoming Interface (srcintf)	wifi
Source Address (srcaddr)	all
Outgoing Interface (dstintf)	port1
Destination Address (dstaddr)	all
Schedule	always
Service	ALL
Action	ACCEPT
Enable NAT	Selected

Controlled AP selection support in FWF client mode

Use the following CLI commands to provide a more controlled AP selection method (supported in FortiWiFi client mode).

Syntax

```
config system interface
  edit {name}
    set wifi-ap-band {any | 5g-preferred | 5g-only}
  next
end
```

Support for location-based services

FortiOS supports location-based services by collecting information about WiFi devices near FortiGate-managed access points, even if the devices don't associate with the network.

Overview

WiFi devices broadcast packets as they search for available networks. The FortiGate WiFi controller can collect information about the interval, duration, and signal strength of these packets. The Euclid Analytics service uses this information to track the movements of the device owner. A typical application of this technology is to analyze shopper behavior in a shopping center. Which stores do people walk past? Which window displays do they stop to look at? Which stores do they enter and how long do they spend there? The shoppers are not personally identified, each is known only by the MAC address of their WiFi device.

After enabling location tracking on the FortiGate unit, you can confirm that the feature is working by using a specialized diagnostic command to view the raw tracking data. The Euclid Analytics service obtains the same data in its proprietary format using a JSON inquiry through the FortiGate unit's web-based manager interface.

Configuring location tracking

You can enable location tracking in any FortiAP profile, using the CLI. Location tracking is part of location-based services. Set the `station-locate` field to `enable`. For example:

```
config wireless-controller wtp-profile
  edit "FAP220B-locate"
    set ap-country US
    config platform
      set type 220B
    end
    config lbs
      set station-locate enable
    end
  end
```

Automatic deletion of outdated presence data

The FortiGate generates a log entry only the first time that station-locate detects a mobile client. No log is generated for clients that have been detected before. To log repeat client visits, previous station presence data must be deleted (flushed). The `sta-locate-timer` can flush this data periodically. The default period is 1800 seconds (30 minutes). The timer can be set to any value between 1 and 86400 seconds (24 hours). A setting of 0 disables the flush, meaning a client is logged only on the very first visit.

The timer is one of the wireless controller timers and it can be set in the CLI. For example:

```
config wireless-controller timers
  set sta-locate-timer 1800
end
```

The `sta-locate-timer` should not be set to less than the `sta-capability-timer` (default 30 seconds) because that could cause duplicate logs to be generated.

FortiPresence push REST API

When the FortiGate is located on a private IP network, the FortiPresence server cannot poll the FortiGate for information. Instead, the FortiGate must be configured to push the information to the FortiPresence server.

Enter the following command:

```
config wireless-controller wtp-profile
  edit "FP223B-GuestWiFi"
    config lbs
      set fortipresence {enable | disable}
      set fortipresence-server <ip-address> Default is 3000.
      set fortipresence-port <port>
      set fortipresence-secret <password>
      set fortipresence-project <name>
      set fortipresence-frequency <5-65535> Default is 30.
      set fortipresence-rogue {enable | disable} Enable/disable reporting of Rogue APs.
      set fortipresence-unassoc {enable | disable} Enable/disable reporting of unassociated devices.
    end
  end
end
```

Viewing device location data on the FortiGate unit

You can use the FortiGate CLI to list located devices. This is mainly useful to confirm that the location data feature is working. You can also reset device location data.

To list located devices

```
diag wireless-controller wlac -c sta-locate
```

To reset device location data

```
diag wireless-controller wlac -c sta-locate-reset
```

Example output

The following output shows data for three WiFi devices.

```
FWF60C3G11004319 # diagnose wireless-controller wlac -c sta-locate
sta_mac vfid rid base_mac freq_lst frm_cnt frm_fst frm_last intv_sum intv2_sum intv3_
sum intv_min intv_max signal_sum signal2_sum signal3_sum sig_min sig_max sig_fst
sig_last ap

00:0b:6b:22:82:61 0
FAP22B3U11005354 0 0 00:09:0f:f1:bb:e4 5745 257 708 56 651 1836 6441 0 12 -21832
1855438 -157758796 -88 -81 -84 -88 0

00:db:df:24:1a:67 0
FAP22B3U11005354 0 0 00:09:0f:f1:bb:e4 5745 42 1666 41 1625 97210 5831613 0 60 -3608
310072 -26658680 -90 -83 -85 -89 0

10:68:3f:50:22:29 0
FAP22B3U11005354 0 0 00:09:0f:f1:bb:e4 5745 102 1623 58 1565 94136 5664566 0 60 -8025
631703 -49751433 -84 -75 -78 -79 0
```

The output for each device appears on two lines. The first line contains only the device MAC address and the VLAN ID. The second line begins with the ID (serial number) of the FortiWiFi or FortiAP unit that detected the device, the AP's MAC address, and then the fields that the Euclid service uses. Because of its length, this line wraps around and displays as multiple lines.

Troubleshooting

In the following section, you will learn basic troubleshooting techniques for a secure Fortinet wireless LAN including:

- strategies for troubleshooting Fortinet wireless devices
- how to avoid common misconfigurations
- solutions to connectivity issues
- capturing and analyzing wireless traffic
- wireless debug commands

The goal of this document is to provide you with practical knowledge that you can use to troubleshoot the FortiOS wireless controller and FortiAP devices. This includes how to use tools and apply CLI commands for maintenance and troubleshooting of your wireless network infrastructure, analyze problems per OSI layer, explore diagnostics for commissioning issues regarding at-client and access point connectivity problems, and understand the packet sniffer technique as a strong troubleshooting tool.

The content is divided as follows:

FortiAP shell command through CAPWAP control tunnel

Very often, the FortiAP in the field is behind a NAT device, and access to the FortiAP through Telnet or SSH is not available. As a troubleshooting enhancement, this feature allows an AP shell command up to 127-bytes sent to the FAP, and FAP will run this command, and return the results to the controller using the CAPWAP tunnel.

The maximum output from a command is limited to 4M, and the default output size is set to 32K.

The FortiAP will only report running results to the controller after the command is finished. If a new command is sent to the AP before the previous command is finished, the previous command will be canceled.

Enter the following:

```
diag w-c wlac wtpcmd wtp_ip wtp_port cmd [cmd-to-ap] cmd: run,show,showhex,clr,r&h,r&sh
```

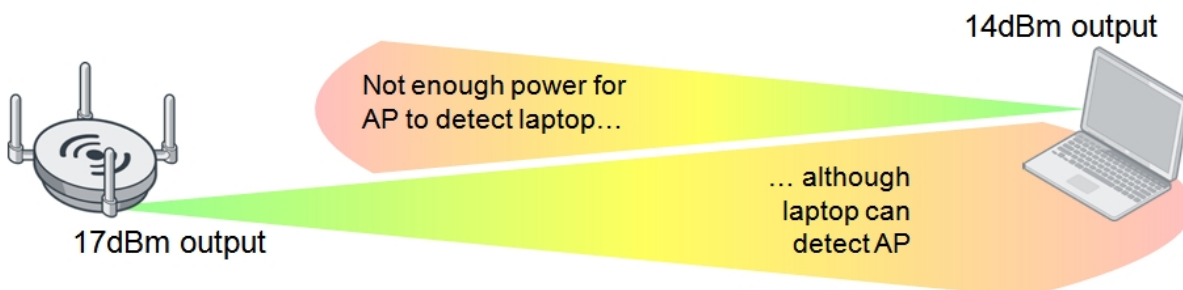
- **cmd-to-ap:** any shell commands, but AP will not report results until the command is finished on the AP
- **run:** controller sends the ap-cmd to the FAP to run
- **show:** show current results reported by the AP in text
- **showhex:** show current results reported by the AP in hex
- **clr:** clear reported results
- **r&s:** run/show
- **r&sh:** run/showhex

Signal strength issues

Poor signal strength is possibly the most common customer complaint. Below you will learn where to begin identifying and troubleshooting poor signal strength, and learn what information you can obtain from the customer to help resolve signal strength issues.

Asymmetric power issue

Asymmetric power issues are a typical problem. Wireless is two-way communication; high power access points (APs) can usually transmit a long distance, however, the client's ability to transmit is usually not equal to that of the AP and, as such, cannot return transmission if the distance is too far.



Measuring signal strength in both directions

To solve an asymmetric power issue, measure the signal strength in both directions. APs usually have enough power to transmit long distances, but sometimes battery-powered clients have a reply signal that has less power, and therefore the AP cannot detect their signal.

It is recommended that you match the transmission power of the AP to the least powerful wireless client—around 10 decibels per milliwatt (dBm) for iPhones and 14dBm for most laptops.

Even if the signal is strong enough, other devices may be emitting radiation as well, causing interference. To identify the difference, read the client Rx strength from the FortiGate GUI (under **Monitor > WiFi Client Monitor**) or CLI.

The **Signal Strength/Noise** value provides the received signal strength indicator (RSSI) of the wireless client. For example, A value of -85dBm to -95dBm is equal to about 10dB levels; this is not a desirable signal strength. In the following screenshot, one of the clients is at 18dB, which is getting close to the perimeter of its range.

SSID	FortiAP	IP	Device	Channel	Bandwidth Tx/Rx	Signal Strength/Noise	Signal
MavisF	FAP28C3X13000119 (1)	10.0.2.8	e8:91:20:90:6e:23	6	1 kbps	29 dB	
MavisF	FP320C3X14000668 (1)	192.168.255.112	1c:69:a5:c8:e8:3e	11	80 bps	35 dB	
MavisF	FP320C3X14000668 (2)	192.168.255.101	58:55:ca:36:28:7d	44	12 kbps	51 dB	
MavisF	FAP28C3X13000119 (1)	10.0.2.9	Acer A1-830 Tablet	6	543 bps	18 dB	
MavisF	FAP28C3X13000119 (1)	10.0.2.13	08:ed:b9:4f:98:ad	6	16 kbps	31 dB	
MavisF	FP320C3X14000668 (1)	192.168.255.115	Ellas_Tablet	11	0 bps	35 dB	



The Signal Strength/Noise value received from the FortiAP by clients, and vice versa, should be within the range of -20dBm to -65dBm.

You can also confirm the transmission (Tx) power of the controller on the AP profile (wtp-profile) and the FortiAP (iwconfig), and check the power management (auto-Tx) options.

Controller configured transmitting power - CLI:

```
config wireless-controller wtp-profile
config <radio>
```

```
show
(the following output is limited to power levels)
auto-power-level : enable
auto-power-high : 17
auto-power-low : 10
```

Actual FortiAP transmitting power - CLI:

```
iwconfig wlan00
```

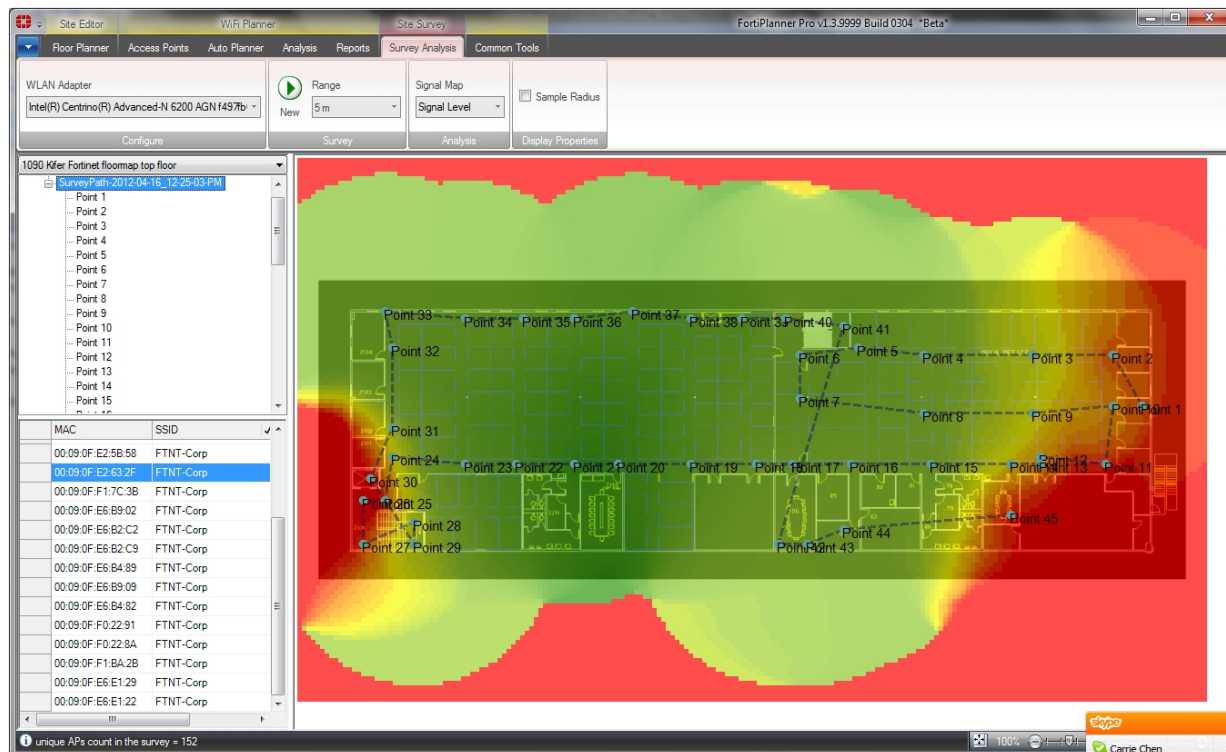
Result:

```
wlan00 IEEE 802.11ng ESSID:"signal-check"
Mode:Master Frequency:2.412 GHz Access Point:<MAC add>
Bit Rate:130 Mb/s Tx-Power=28 dBm
```

Using FortiPlanner PRO with a site survey

The most thorough method to solve signal strength issues is to perform a site survey. To this end, Fortinet offers the FortiPlanner, downloadable at http://www.fortinet.com/resource_center/product_downloads.html.

Sample depiction of a site survey using FortiPlanner



The site survey provides you with optimal placement for your APs based on the variables in your environment. You must provide the site survey detailed information including a floor plan (to scale), structural materials, and more. It will allow you to place the APs on the map and adjust the radio bands and power levels while providing you with visual wireless coverage.

Below is a list of mechanisms for gathering further information on the client for Rx strength. The goal is to see how well the client is receiving the signal from the AP. You can also verify FortiAP signal strength on the client

using WiFi client utilities, or third party utilities such as InSSIDer or MetaGeek Chanalyzer. You can get similar tools from the app stores on Android and iOS devices.

- Professional Site Survey software (Ekahau, Airmagnet survey Pro, FortiPlanner)
- InSSIDer
- On Windows: “*netsh wlan show networks mode=bssid*” (look for the BSSID, it's in % not in dBm!)
- On MacOS: Use the “*airport*” command:

```
"/System/Library/PrivateFrameworks/Apple80211.framework/Versions/A/Resources/airport" airport -s | grep <the_bssid>
```

 (live scan each time)
- On Droid: WiFiFoFum

Frequency interference

If the wireless signal seems to be strong but then periodically drops, this may be a symptom of frequency interference. Frequency interference is when another device also emits radio frequency using the same channel, co-channel, or adjacent channel, thereby overpowering or corrupting your signal. This is a common problem on a 2.4GHz network.

There are two types of interference: coherent and non-coherent.

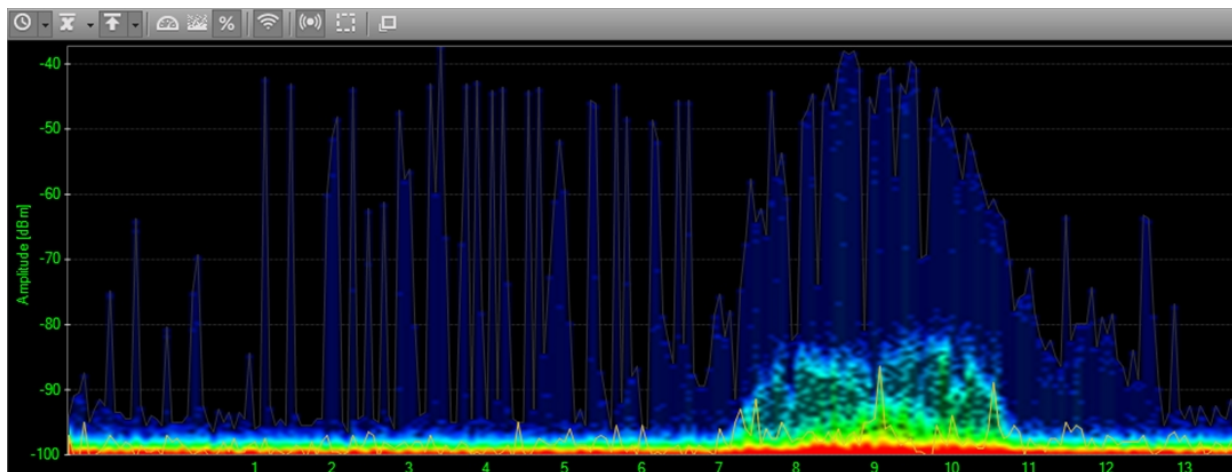
- **Coherent interference:** a result of another device using the same channel as your AP, or poor planning of a wireless infrastructure (perhaps the other nearby APs are using the same channel or the signal strength is too high).
- **Non-coherent interference:** a result of other radio signals such as bluetooth, microwave, cordless phone, or (as in medical environments) x-ray machines.

Most common and simple solution for frequency interference is to change your operation channel. Typically, the channel can be set from 1 to 11 for the broadcast frequency, although you should always use channels 1, 6, and 11 on the 2.4GHz band.

Another solution, if it's appropriate for your location, is to use the 5GHz band instead.

MetaGeek Chanalyzer

You can perform a site survey using spectrum analysis at various points in your environment looking for signal versus interference/noise. MetaGeek Chanalyzer is an example of a third party utility which shows a noise threshold.



Note that a signal of -95dBm or less will be ignored by Fortinet wireless adapters.

Throughput issues

Sometimes communication issues can be caused by low performance.

Testing the link

You can identify delays or lost packets by sending ping packets from your wireless client. If there is more than 10ms of delay, there may be a problem with your wireless deployment, such as:

- a weak transmit signal from the client (the host does not reach the AP)
- the AP utilization is too high (your AP could be saturated with connected clients)
- interference (third party signal could degrade your AP or client's ability to detect signals between them)
- weak transmit power from the AP (the AP does not reach the host) -- not common in a properly deployed network, unless the client is too far away

Keep in mind that water will also cause a reduction in radio signal strength for those making use out of outdoor APs or wireless on a boat.

Performance testing


If the FortiAP gives bad throughput to the client, the link may drop. The throughput or performance can be measured on your smartphone with third party applications tool such as iPerf and jPerf.

Measuring file transfer speed

Another way to get a sense of your throughput issues is to measure the speed of a file transfer on your network. Create a test file at a specific size and measure the speed at which Windows measures the transfer. The command below will create a 50MB file.

- `fsutil file createnew test.txt 52428800`

The following image shows a network transfer speed of just over 24Mbps. The theoretical speed of 802.11g is 54Mbps, which is what this client is using. A wireless client is never likely to see the theoretical speed.



TKIP limitation

If you find that throughput is a problem, avoid WPA security encrypted with Temporal Key Integrity Protocol (TKIP) as it supports communications only at 54Mbps. Use WPA-2 AES instead.

Speeds are very much based on what the client computer can handle as well. The maximum client connection rate of 130Mbps is for 2.4GHz on a 2x2, or 300Mbps for 5Ghz on a 2x2 (using shortguard and channel bonding enabled).

If you want to get more than 54Mbps with 802.11n, do not use legacy TKIP, use CCMP instead. This is standard for legacy compatibility.

Preventing IP fragmentation in CAPWAP

TKIP is not the only possible source of decreased throughput. When a wireless client sends jumbo frames using a CAPWAP tunnel, it can result in data loss, jitter, and decreased throughput.

Using the following commands you can customize the uplink rates and downlink rates in the CAPWAP tunnel to prevent fragmentation and avoid data loss.

```
config wireless-controller wtp
  edit new-wtp
    set override-ip-fragment enable
    set ip-fragment-preventing [tcp-mss-adjust | icmp-unreachable]
    set tun-mtu-uplink [0 | 576 | 1500]
    set tun-mtu-downlink [0 | 576 | 1500]
  end
end
```

The default value is 0, however the recommended value will depend on the type of traffic. For example, IPsec in tunnel mode has 52 bytes of overhead, so you might use 1400 or less for uplink and downlink.

Slowness in the DTLS response

It's important to know all the elements involved in the CAPWAP association:

- Request
- Response
- DTLS
- Join
- Configuration

All of these are bidirectional. So if the DTLS response is slow, this might be the result of a configuration error. This issue can also be caused by a certificate during discovery response. You can read more about this in [RFC 5416](#).

Connection issues

If the client has a connectivity issue that is not due to signal strength, the solution varies by the symptom.

Client connection issues

1. If client is unable to connect to FortiAP:

- Make sure the client's security and authentication settings match with FortiAP and check the certificates as well.
 - Try upgrading the Wi-Fi adapter driver and FortiGate/FortiAP firmware.
 - If other clients can connect, it could be interoperability; run debug commands and sniffer packets.
 - Look for rogue suppression by sniffing the wireless traffic and looking for the disconnect in the output (using the AP or wireless packet sniffer).
 - Try changing the IEEE protocol from 802.11n to 802.11bg or 802.11a only.
2. If the client drops and reconnects:
 - The client might be de-authenticating periodically. Check the sleep mode on the client.
 - The issue could be related to power-saver settings. The client may need to update drivers.
 - The issue could also be caused by flapping between APs. Check the roaming sensitivity settings on the client or the preferred wireless network settings on the client—if another WiFi network is available, the client may connect to it if it is a preferred network. Also, check the DHCP configuration as it may be an IP conflict.
 3. If the client drops and never connects:
 - It could have roamed to another SSID, so check the standby and sleep modes.
 - You may need to bring the interface up and down.
 4. If the client connects, but no IP address is acquired by the client:
 - Check the DHCP configuration and the network.
 - It could be a broadcast issue, so check the WEP encryption key and set a static IP address and VLANs.

Debug

You should also enable client debug on the controller for problematic clients to see the stage at which the client fails to connect. Try to connect from the problematic client and run the following debug command, which allows you to see the four-way handshake of the client association:

```
diagnose wireless-controller wlac sta_filter <client MAC address> 2
```

Example of a successful client connection:

The following is a sample debug output for the above command, with successful association/DHCP phases and PSK key exchange (identified in color):

```
FG600B3909600253 #
91155.197 <ih> IEEE 802.11 mgmt::assoc_req <== 30:46:9a:f9:fa:34 vap signal-check rId 0
wId 0 00:09:0f:f3:20:45
91155.197 <ih> IEEE 802.11 mgmt::assoc_resp ==> 30:46:9a:f9:fa:34 vap signal-check rId 0
wId 0 00:09:0f:f3:20:45 resp 0
91155.197 <cc> STA_CFG_REQ(15) sta 30:46:9a:f9:fa:34 add ==> ws (0-192.168.35.1:5246) rId 0
wId 0
91155.197 <dc> STA add 30:46:9a:f9:fa:34 vap signal-check ws (0-192.168.35.1:5246) rId 0
wId 0 bssid 00:09:0f:f3:20:45 NON-AUTH
91155.197 <cc> STA add 30:46:9a:f9:fa:34 vap signal-check ws (0-192.168.35.1:5246) rId 0
wId 0 00:09:0f:f3:20:45 sec WPA2 AUTO auth 0
91155.199 <cc> STA_CFG_RESP(15) 30:46:9a:f9:fa:34 <== ws (0-192.168.35.1:5246) rc 0
(Success)
91155.199 <eh> send 1/4 msg of 4-Way Handshake
91155.199 <eh> send IEEE 802.1X ver=1 type=3 (EAPOL_KEY) data len=95 replay cnt 1
91155.199 <eh> IEEE 802.1X (EAPOL 99B) ==> 30:46:9a:f9:fa:34 ws (0-192.168.35.1:5246) rId
0 wId 0 00:09:0f:f3:20:45
91155.217 <eh> IEEE 802.1X (EAPOL 121B) <== 30:46:9a:f9:fa:34 ws (0-192.168.35.1:5246) rId
0 wId 0 00:09:0f:f3:20:45
```

```

91155.217 <eh> recv IEEE 802.1X ver=1 type=3 (EAPOL_KEY) data len=117
91155.217 <eh> recv EAPOL-Key 2/4 Pairwise replay cnt 1
91155.218 <eh> send 3/4 msg of 4-Way Handshake
91155.218 <eh> send IEEE 802.1X ver=1 type=3 (EAPOL_KEY) data len=175 replay cnt 2
91155.218 <eh> IEEE 802.1X (EAPOL 179B) ==> 30:46:9a:f9:fa:34 ws (0-192.168.35.1:5246) rId
0 wId 0 00:09:0f:f3:20:45
91155.223 <eh> IEEE 802.1X (EAPOL 99B) <== 30:46:9a:f9:fa:34 ws (0-192.168.35.1:5246) rId
0 wId 0 00:09:0f:f3:20:45
91155.223 <eh> recv IEEE 802.1X ver=1 type=3 (EAPOL_KEY) data len=95
91155.223 <eh> recv EAPOL-Key 4/4 Pairwise replay cnt 2
91155.223 <dc> STA chg 30:46:9a:f9:fa:34 vap signal-check ws (0-192.168.35.1:5246) rId 0
wId 0 bssid 00:09:0f:f3:20:45 AUTH
91155.224 <cc> STA chg 30:46:9a:f9:fa:34 vap signal-check ws (0-192.168.35.1:5246) rId 0
wId 0 00:09:0f:f3:20:45 sec WPA2 AUTO auth 1
91155.224 <cc> STA_CFG_REQ(16) sta 30:46:9a:f9:fa:34 add key (len=16) ==> ws (0-
192.168.35.1:5246) rId 0 wId 0
91155.226 <cc> STA_CFG_RESP(16) 30:46:9a:f9:fa:34 <== ws (0-192.168.35.1:5246) rc 0
(Success)
91155.226 <eh> ***pairwise key handshake completed*** (RSN)
91155.257 <dc> DHCP Request server 0.0.0.0 <== host ADMINFO-FD4I2HK mac 30:46:9a:f9:fa:34
ip 172.16.1.16
91155.258 <dc> DHCP Ack server 172.16.1.1 ==> host mac 30:46:9a:f9:fa:34 ip 172.16.1.16
mask 255.255.255.0 gw 172.16.1.1

```

where:

- **orange** represents the association phase,
- **blue** represents the PSK exchange,
- and **green** represents the DHCP phase.

It is important to note the messages for a correct association phase, four-way handshake, and DHCP phase.

Checking WiFi password

Admins can view plain text passwords (captive-portal-radius-secret and passphrase) under config wireless-controller vap.

Note that security must be set as a WPA-personal setting.

FortiAP connection issues

Clients are not the only device that can fail to connect, of course. A communication problem could arise from the FortiAP.

Some examples include:

- The FortiAP is not connecting to the wireless controller.
- One FortiAP intermittently disconnects and re-connects.
- All FortiAPs intermittently disconnect and re-connect.
- Unable to Telnet to FortiAP from controller/administrator workstation.

In the above cases:

- Check networking on the distribution system for all related FortiAPs.
- Check the authorization status of managed APs from the wireless controller.

- Restart the `cw_acd` process (**Note:** All APs will drop if you do this, and you may be troubleshooting just one AP).
- Check the controller crash log for any wireless controller daemon crash using the following command:

```
diagnose debug crashlog read
```

Debug

For a quick assessment of the association communication between the controller and the FortiAP, run the following sniffer command to see if you can verify that the AP is communicating to the controller by identifying the CAPWAP communication:

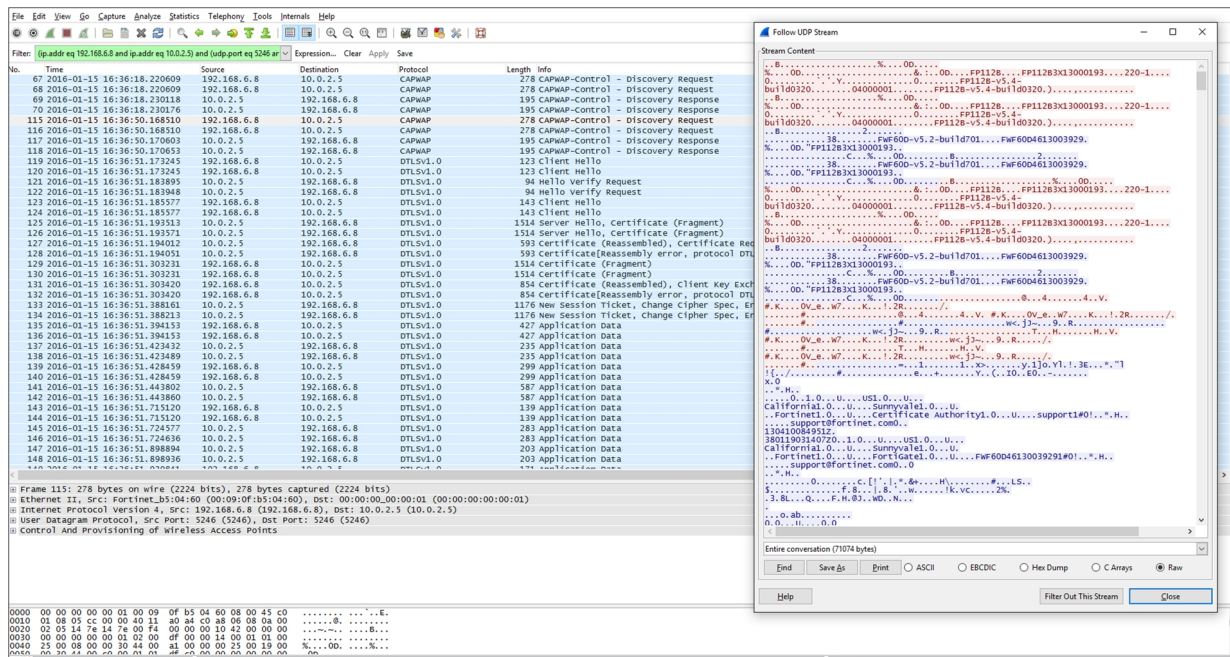
```
diagnose sniff packet <interface_name> "port 5246" 4
```

If you do not see this communication, then you can investigate the network or the settings on the AP to see why it is not reaching the controller.

The following command allows you to collect verbose output from the sniff that can be converted to a PCAP and viewed in Wireshark.

```
diagnose sniff packet <interface_name> "port 5246" 6 o 1
```

The image below shows the beginning of the AP's association to the controller. You can see the discovery Request and Response at the top.



Throughout debugging it is recommended to:

- Enable Telnet login to the FortiAP device so that you can log in and issue local debugging commands:

```
config wireless-controller wtp
edit "<FortiAP_serial_number>"
set override-allowaccess {disable|enable}
set allowaccess {telnet | http | https | ssh}
end
```

- Try to connect to the wireless controller from the problematic FortiAP to verify routes exist.

- Enable wtp (FortiAP) debugging on the wireless controller for problematic FortiAPs to determine the point at which the FortiAP fails to connect:

```
diag wireless-controller wlap wtp_filter FP112B3X13000193 0-192.168.6.8:5246 2
(replace the serial number and IP address of the FortiAP)
di de console timestamp en
di de application cw_acd 0x7ff
di de en
```

Example of a successful AP and controller association:

The previous debug command provides similar output to the sample debug message below for a successful association between the FortiAP and the wireless controller. This includes the elements of the CAPWAP protocol; the Request, Response, DTLS, Join, and Configuration (identified in color). All of these are bi-directional, so if the DTLS response is slow, it may be an example of a configuration error.

```
56704.575 <msg> DISCOVERY_REQ (12) <== ws (0-192.168.35.1:5246)
56704.575 <msg> DISCOVERY_RESP (12) ==> ws (0-192.168.35.1:5246)
56707.575 <msg> DISCOVERY_REQ (13) <== ws (0-192.168.35.1:5246)
56707.575 <msg> DISCOVERY_RESP (13) ==> ws (0-192.168.35.1:5246)
56709.577 <aev> - CWAE_INIT_COMPLETE ws (0-192.168.35.1:5246)
56709.577 <aev> - CWAE_LISTENER_THREAD_READY ws (0-192.168.35.1:5246)
56709.577 <fsm> old CWAS_START(0) ev CWAE_INIT_COMPLETE(0) new CWAS_IDLE(1)
56709.577 <fsm> old CWAS_IDLE(1) ev CWAE_LISTENER_THREAD_READY(1) new CWAS_DTLS_SETUP(4)
56709.623 <aev> - CWAE_DTLS_PEER_ID_RECV ws (0-192.168.35.1:5246)
56709.623 <aev> - CWAE_DTLS_AUTH_PASS ws (0-192.168.35.1:5246)
56709.623 <aev> - CWAE_DTLS_ESTABLISHED ws (0-192.168.35.1:5246)
56709.623 <fsm> old CWAS_DTLS_SETUP(4) ev CWAE_DTLS_PEER_ID_RECV(7) new CWAS_DTLS_
AUTHORIZE(2)
56709.623 <fsm> old CWAS_DTLS_AUTHORIZE(2) ev CWAE_DTLS_AUTH_PASS(3) new CWAS_DTLS_CONN(5)
56709.623 <fsm> old CWAS_DTLS_CONN(5) ev CWAE_DTLS_ESTABLISHED(8) new CWAS_JOIN(7)
56709.625 <msg> JOIN_REQ (14) <== ws (0-192.168.35.1:5246)
56709.625 <aev> - CWAE_JOIN_REQ_RECV ws (0-192.168.35.1:5246)
56709.626 <fsm> old CWAS_JOIN(7) ev CWAE_JOIN_REQ_RECV(12) new CWAS_JOIN(7)
56709.629 <msg> CFG_STATUS (15) <== ws (0-192.168.35.1:5246)
56709.629 <aev> - CWAE_CFG_STATUS_REQ ws (0-192.168.35.1:5246)
56709.629 <fsm> old CWAS_JOIN(7) ev CWAE_CFG_STATUS_REQ(13) new CWAS_CONFIG(8)
56710.178 <msg> CHG_STATE_EVENT_REQ (16) <== ws (0-192.168.35.1:5246)
56710.178 <aev> - CWAE_CHG_STATE_EVENT_REQ_RECV ws (0-192.168.35.1:5246)
56710.178 <fsm> old CWAS_CONFIG(8) ev CWAE_CHG_STATE_EVENT_REQ_RECV(23) new CWAS_DATA_
CHAN_SETUP(10)
56710.220 <aev> - CWAE_DATA_CHAN_CONNECTED ws (0-192.168.35.1:5246)
56710.220 <msg> DATA_CHAN_KEEP_ALIVE <== ws (0-192.168.35.1:5246)
56710.220 <aev> - CWAE_DATA_CHAN_KEEP_ALIVE_RECV ws (0-192.168.35.1:5246)
56710.220 <msg> DATA_CHAN_KEEP_ALIVE ==> ws (0-192.168.35.1:5246)
56710.220 <fsm> old CWAS_DATA_CHAN_SETUP(10) ev CWAE_DATA_CHAN_CONNECTED(32) new CWAS_
DATA_CHECK(11)
56710.220 <aev> - CWAE_DATA_CHAN_VERIFIED ws (0-192.168.35.1:5246)
56710.220 <fsm> old CWAS_DATA_CHECK(11) ev CWAE_DATA_CHAN_KEEP_ALIVE_RECV(35) new CWAS_
DATA_CHECK(11)
56710.220 <fsm> old CWAS_DATA_CHECK(11) ev CWAE_DATA_CHAN_VERIFIED(36) new CWAS_RUN(12)
56710.228 <msg> WTP_EVENT_REQ (17) <== ws (0-192.168.35.1:5246)
56710.228 <aev> - CWAE_WTP_EVENT_REQ_RECV ws (0-192.168.35.1:5246)
56710.228 <fsm> old CWAS_RUN(12) ev CWAE_WTP_EVENT_REQ_RECV(42) new CWAS_RUN(12)
56710.230 <msg> CFG_UPDATE_RESP (1) <== ws (0-192.168.35.1:5246) rc 0 (Success)
56710.230 <aev> - CWAE_CFG_UPDATE_RESP_RECV ws (0-192.168.35.1:5246)
56710.230 <msg> WTP_EVENT_REQ (18) <== ws (0-192.168.35.1:5246)
56710.230 <aev> - CWAE_WTP_EVENT_REQ_RECV ws (0-192.168.35.1:5246)
```

```

56710.230 <fsm> old CWAS_RUN(12) ev CWAE_CFG_UPDATE_RESP_RECV(37) new CWAS_RUN(12)
56710.230 <fsm> old CWAS_RUN(12) ev CWAE_WTP_EVENT_REQ_RECV(42) new CWAS_RUN(12)
56710.231 <msg> WTP_EVENT_REQ (19) <== ws (0-192.168.35.1:5246)
56710.231 <aev> - CWAE_WTP_EVENT_REQ_RECV ws (0-192.168.35.1:5246)
56710.231 <fsm> old CWAS_RUN(12) ev CWAE_WTP_EVENT_REQ_RECV(42) new CWAS_RUN(12)
56710.232 <msg> CFG_UPDATE_RESP (2) <== ws (0-192.168.35.1:5246) rc 0 (Success)
56710.232 <aev> - CWAE_CFG_UPDATE_RESP_RECV ws (0-192.168.35.1:5246)
56710.232 <fsm> old CWAS_RUN(12) ev CWAE_CFG_UPDATE_RESP_RECV(37) new CWAS_RUN(12)
56710.233 <msg> WTP_EVENT_REQ (20) <== ws (0-192.168.35.1:5246)
56710.233 <aev> - CWAE_WTP_EVENT_REQ_RECV ws (0-192.168.35.1:5246)
56710.233 <fsm> old CWAS_RUN(12) ev CWAE_WTP_EVENT_REQ_RECV(42) new CWAS_RUN(12)
56712.253 < . > AC (2) -> WTP (0-192.168.35.1:5246) State: CWAS_RUN (12) accept 3 live 3
           dbg 00000000 pkts 12493 0
56715.253 < . > AC (2) -> WTP (0-192.168.35.1:5246) State: CWAS_RUN (12) accept 3 live 6
           dbg 00000000 pkts 12493 0
56718.253 < . > AC (2) -> WTP (0-192.168.35.1:5246) State: CWAS_RUN (12) accept 3 live 9
           dbg 00000000 pkts 12493 0
56719.253 <aev> - CWAE_AC_ECHO_INTV_TMR_EXPIRE ws (0-192.168.35.1:5246)
56719.253 <fsm> old CWAS_RUN(12) ev CWAE_AC_ECHO_INTV_TMR_EXPIRE(39) new CWAS_RUN(12)
56719.576 <msg> ECHO_REQ (21) <== ws (0-192.168.35.1:5246)
56719.576 <aev> - CWAE_ECHO_REQ_RECV ws (0-192.168.35.1:5246)
56719.577 <fsm> old CWAS_RUN(12) ev CWAE_ECHO_REQ_RECV(27) new CWAS_RUN(12)

```

where:

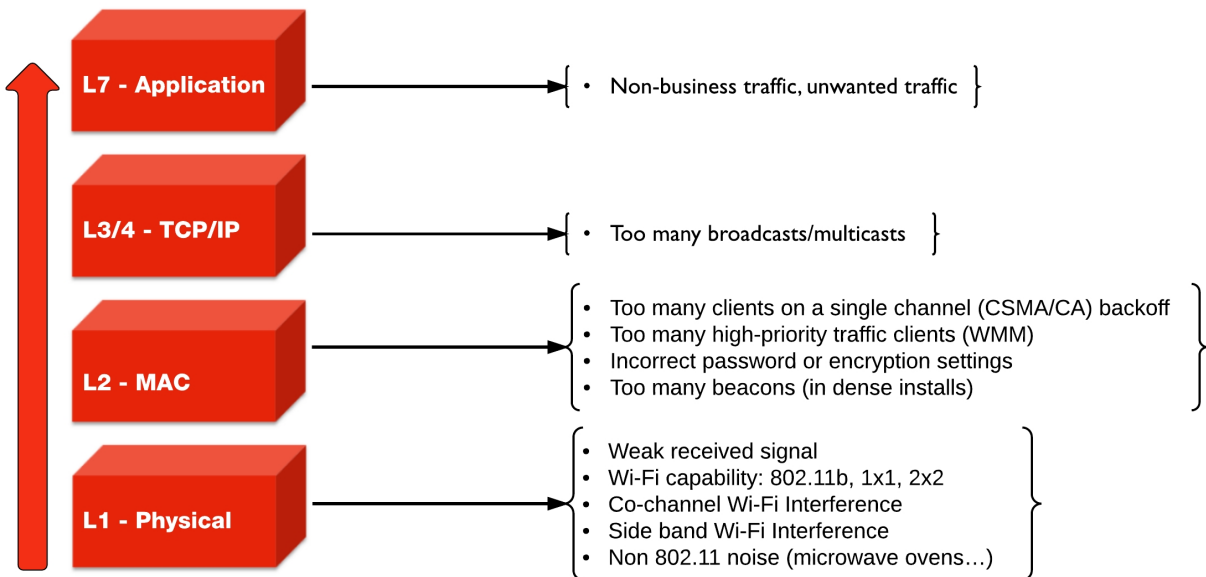
- **orange** represents the Discovery phase,
- **blue** indicates that the control channels have been established using DTLS,
- **green** represents the access point Discovery and Join phase,
- **purple** represents the Clear Text channel,
- and **pink** indicates that the FortiAP successfully connected to the wireless controller.

General problems

Not all WiFi problems are related to signal strength, interference, or misconfiguration. The following OSI model identifies some of the more common issues per layer.

Best practices for troubleshooting vary depending on the affected layer (see below).

Common sources of wireless issues



Best practices for Layer 1

Common physical layer issues include:

- Weak received signal,
- WiFi capability: 802.11b, 1x1, 2x2,
- Co-channel WiFi interference,
- Side band WiFi interference,
- Non 802.11 noise (microwave ovens...).

To avoid physical layer issues:

- Determine RST (Receiver Sensitivity Threshold) for your device, or use -70dBm as a rule of thumb.
- Match AP TX output power to the client TX output power.
 - **Note:** iPhone TX power is only 10dBm.
- Use DFS (Dynamic Frequency Selection) for high performance data 20/40 MHz.
- Use 5GHz UNII-1 & 3 (Non-DFS) bands with static channel assignment for latency-sensitive applications.
- Do not use 40MHz channels in 2.4 GHz band (channel bonding is not allowed in FortiOS).

Best practices for Layer 2

Common data link (MAC) layer issues include:

- Too many clients on a single channel (CSMA/CA) backoff,
- Too many high-priority traffic clients (WMM),
- Incorrect password or encryption settings,
- Too many beacons (in dense installs).

To avoid data link layer issues:


- Only use CCMP/AES (WPA2) encryption (not TKIP).
- In high density deployments, turn off SSID broadcast or turn down SSID rates. Review and possibly reduce the beacon interval.
- Determine the best cell size for applications:
 - For few users and low bandwidth latency sensitive applications, use high transmit power to create larger cells.
 - For high performance/high capacity installations, use lower transmit power to create smaller cells (set FortiPlanner at 10dBm TX power), but bear in mind that this will require more roaming.

Cells and co-channel interference

In high density deployments, multiple APs are used, and each one services an area called a cell. However, these cells can cause interference with each other. This is a common problem. The radio signal from one AP interferes with, or cancels out, the radio signal from another AP.


In the following diagram, note the interference zone created by one radio, causing interference on its neighbouring APs.

The interference zone can be twice the radius of the signal, and the signal at its edge can be -67dBm.



Reducing co-channel interference

For best results, use a 'honeycomb' pattern as a deployment strategy. The idea is to *stagger* repeated channels furthest from each other to avoid interference.



Best practices for Layer 3 and above

For TCP/IP layers and above, a common source of latency, or slowness in the wireless traffic, is too many broadcasts or multicasts. These types of issues can result from non-business and/or unwanted traffic.

To resolve issues at the TCP/IP layer and above:

- Identify business-critical applications.
- Use Application Control, Web Filtering, Traffic Shaping, and QoS to prioritize applications.
 - Identify unwanted traffic, high-bandwidth web-related traffic, and use Security Profiles.
 - Use the traffic shaper on a policy to rate-limit this traffic.

These configurations are performed directly on the FortiGate.

Packet sniffer

Capturing the traffic between the controller and the FortiAP can help you identify most FortiAP and client connection issues.

This section describes the following recommended packet sniffing techniques:

- [CAPWAP packet sniffer](#)
- [Wireless traffic packet sniffer](#)

CAPWAP packet sniffer

The first recommended technique consists of sniffing the CAPWAP traffic.

- Enable plain control on the controller and on the FortiAP to capture clear control traffic on UDP port 5246.
 - On the controller:

```
diagnose wireless-controller wlac plain-ctl <FortiAP_serial_number> 1
```

Result:

```
WTP 0-FortiAP2223X11000107 Plain Control: enabled
```

- On the FortiAP:

```
cw_diag plain-ctl 1
```

Result:

```
Current Plain Control: enabled
```

Note that some issues are related to the keep-alive for control and data channel.

- Data traffic on UDP port 5247 is not encrypted. The data itself is encrypted by the wireless security mechanism.

Data traffic is helpful to troubleshoot most of the issues related to station association, EAP authentication, WPA key exchange, roaming, and FortiAP configuration.

You can also set up a host or server to which you can forward the CAPWAP traffic:

1. Configure the host/server to which CAPWAP traffic is forwarded:

```
diagnose wireless-controller wlac sniff-cfg <Host_IP_address> 88888
```

Result:

```
Current Sniff Server: 192.168.25.41, 23352
```

2. Choose which traffic to capture, the interface to which the FortiAP is connected, and the FortiAP's serial number:

```
diagnose wireless-controller wlac sniff <interface_name> <FortiAP_serial_number> 2
```

Result:

```
WTP 0-FortiAP2223X11000107 Sniff: intf port2 enabled (control and data message)
```

In the above syntax, the '2' captures the control and data message—'1' would capture only the control message, and '0' would disable it.

3. Run Wireshark on the host/server to capture CAPWAP traffic from the controller.
 - Decode the traffic as IP to check inner CAPWAP traffic.

Example CAPWAP packet capture

The following image shows an example of a CAPWAP packet capture, where you can see: the Layer 2 header; the sniffed traffic encapsulated into Internet Protocol for transport; CAPWAP encapsulated into UDP for sniffer purpose and encapsulated into IP; CAPWAP control traffic on UDP port 5246; and CAPWAP payload.

No.	Time	Source	Destination	Protocol	Info
1	0.000000	192.168.35.82	192.168.35.80	CAPWAP	Control Msg - Echo Request
2	0.000308	192.168.35.82	192.168.35.80	CAPWAP	Control Msg - Echo Request
3	0.000452	192.168.35.80	192.168.35.82	CAPWAP	Control Msg - Echo Response
4	0.000454	192.168.35.80	192.168.35.82	CAPWAP	Control Msg - Echo Response

Frame 4: 134 bytes on wire (1072 bits), 134 bytes captured (1072 bits)

Ethernet II, Src: Fortinet_c5:ce:66 (00:09:0f:c5:ce:66), Dst: Intel_0e:e3:79 (00:07:e9:0e:e3:79)

Internet Protocol, Src: 192.168.35.80 (192.168.35.80), Dst: 192.168.35.45 (192.168.35.45)

User Datagram Protocol, Src Port: 8887 (8887), Dst Port: 55555 (55555)

Internet Protocol, Src: 192.168.35.80 (192.168.35.80), Dst: 192.168.35.82 (192.168.35.82)

Version: 4
Header length: 20 bytes
Differentiated Services Field: 0x00 (DSCP 0x00: Default; ECN: 0x00)
Total Length: 92
Identification: 0x0021 (33)
Flags: 0x00
Fragment offset: 0
Time to live: 64
Protocol: UDP (17)
Header checksum: 0xb27d [correct]
Source: 192.168.35.80 (192.168.35.80)
Destination: 192.168.35.82 (192.168.35.82)

User Datagram Protocol, Src Port: capwap-control (5246), Dst Port: capwap-control (5246)

CAPWAP Protocol
CAPWAP Header 8 bytes
CAPWAP Control Header 8 bytes
CAPWAP Message Elements 48 bytes

0000 00 07 e9 0e e3 79 00 09 0f c5 ce 66 08 00 45 00y...f..E..
0010 00 78 b1 b7 00 00 40 11 00 f0 c0 a8 23 50 c0 a8 ..x...@...#P..
0020 23 2d 22 b7 d9 03 00 64 00 00 45 00 00 5c 00 21 #-"...d...E..!
0030 00 00 40 11 b2 7d c0 a8 23 50 c0 a8 23 52 14 7e ..@...}...#P..#R..
0040 14 7e 00 48 00 00 00 10 42 00 00 00 00 00 00 00 ~.H....B.....
0050 00 0e 8d 00 33 00 00 25 00 2c 00 00 30 44 00 323...%...OD.2

Frame (frame), 134 bytes Packets: 4 Displayed: 4 Marked: 0 Load time: 0:00.218 Profile: Default

Wireless traffic packet sniffer

The second recommended technique consists of sniffing the wireless traffic directly 'on the air' using your FortiAP.

Wireless traffic packet capture

Packet captures are useful for troubleshooting all wireless client related issues because you can verify data rate and 802.11 parameters, such as radio capabilities, and determine issues with wireless signal strength, interference, or congestion on the network.

A radio can only capture one frequency at a time; one of the radios is set to sniffer mode depending on the traffic or channel required. You must use two FortiAPs to capture both frequencies at the same time.

- Set a radio on the FortiAP to monitor mode.

```
iwconfig wlan10
```

Result:

```
wlan10 IEEE 802.11na ESSID:""  
Mode:Monitor Frequency:5.18 GHz Access Point: Not-Associated
```

- The capture file is stored under the temp directory as *wl_sniff.pcap*

```
/tmp/wl_sniff.pcap
```

- Remember that the capture file is only stored temporarily. If you want to save it, upload it to a TFTP server before rebooting or changing the radio settings.
- The command `cp wl_sniff.pcap newname.pcap` allows you to rename the file.

- Rather than TFTP the file, you can also log in to the AP and retrieve the file via the web interface. Move the file using the command: `mv name /usr/www`

You can verify the file was moved using the command `cd/usr/www` and then browsing to: `<fortiAP_IP>/filename`

Syntax

The following syntax demonstrates how to set the radio to sniffer mode (configurable from the CLI only). Sniffer mode provides options to filter for specific traffic to capture. Notice that you can determine the buffer size, which channel to sniff, the AP's MAC address, and select if you want to sniff the beacons, probes, controls, and data channels.

```
configure wireless-controller wtp-profile
  edit <profile_name>
    configure <radio>
      set mode sniffer
      set ap-sniffer-bufsize 32
      set ap-sniffer-chan 1
      set ap-sniffer-addr 00:00:00:00:00:00
      set ap-sniffer-mgmt-beacon enable
      set ap-sniffer-mgmt-probe enable
      set ap-sniffer-mgmt-other enable
      set ap-sniffer-ctl enable
      set ap-sniffer-data enable
    end
  end
```

Once you've performed the previous CLI configuration, you'll be able to see the packet sniffer mode selected in the GUI dashboard under **WiFi & Switch Controller > FortiAP Profiles** and **WiFi & Switch Controller > Managed FortiAPs**. Bear in mind that if you change the mode from the GUI, you'll have to return to the CLI to re-enable the Sniffer mode.

To disable the sniffer profile in the CLI, use the following commands:

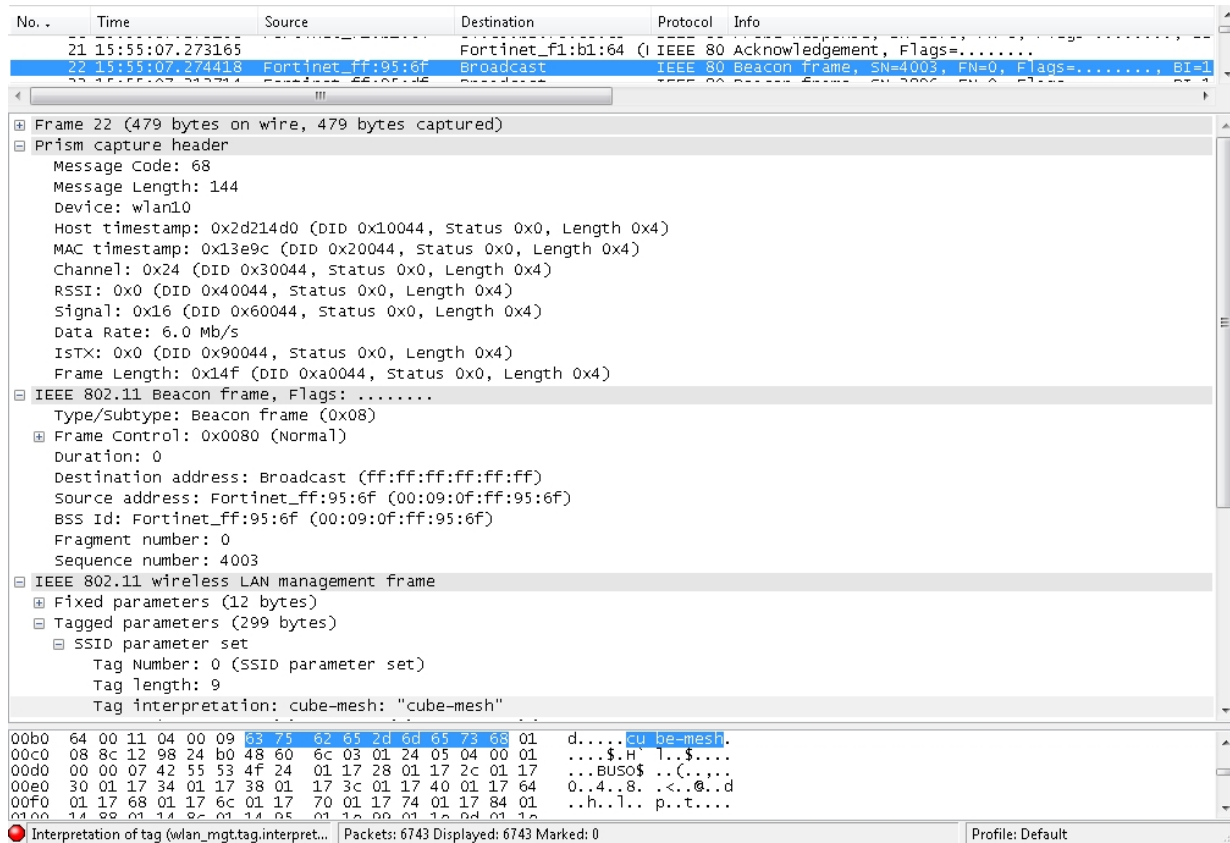
```
config wireless-controller wtp-profile
  edit <profile_name>
    config <radio>
      set ap-sniffer-mgmt-beacon disable
      set ap-sniffer-mgmt-probe disable
      set ap-sniffer-mgmt-other disable
      set ap-sniffer-ctl disable
      set ap-sniffer-data disable
    end
  end
```



If you change the radio mode before sending the file *wl_sniff.cap* to an external TFTP, the file will be deleted and you will lose your packet capture.

Example AP packet capture

The following image shows an example of the AP packet capture. Note the capture header showing channel 36; the beacon frame; the source, destination, and BSSID of the beacon frame; and the SSID of the beacon frame.



Useful debugging commands

For a comprehensive list of useful debug options you can use the following help commands on the controller:

```
diagnose wireless-controller wlac help
(this command lists the options available that pertain to the wireless controller)
```

```
diagnose wireless-controller wlwtp help
(this command lists the options available that pertain to the AP)
```

Sample outputs

Syntax

```
diagnose wireless-controller wlac -c vap
(this command lists the information about the virtual access point, including its MAC address, the BSSID, its SSID, the interface name, and the IP address of the APs that are broadcasting it)
```

Result:

```
bssid          ssid   intf    vfid:ip-port  rId  wId
00:09:0f:d6:cb:12 Office Office  ws (0-192.168.3.33:5246) 0 0
00:09:0f:e6:6b:12 Office Office  ws (0-192.168.1.61:5246) 0 0
06:0e:8e:27:dc:48 Office Office  ws (0-192.168.3.36:5246) 0 0
0a:09:0f:d6:cb:12 public publicAP ws (0-192.168.3.33:5246) 0 1
```

Syntax

```
diagnose wireless-controller wlac -c darrp
```

(this command lists the information pertaining to the radio resource provisioning statistics, including the AP serial number, the number of channels set to choose from, and the operation channel. Note that the 5GHz band is not available on these APs listed)

Result:

wtp_id	rId	base_mac	index	nr_chan	vfid	5G	oper_chan	age
FAP22A3U10600400	0	00:09:0f:d6:cb:12	0	3	0	No	1	87588
FW80CM3910601176	0	06:0e:8e:27:dc:48	1	3	0	No	6	822

Support for extension information for wtp, vap, and station

You can enable or disable extension information at `wtp-profile`, and use the diagnose option below to print out the detail of extension information.

Syntax

```
config wireless-controller wtp-profile
edit test
set lldp [enable | disable]
set ext-info [enable | disable] --> Enable/disable station/VAP/radio extension information.
end
end

diagnose wireless-controller wlac -d [wtp | vap | sta]
```

where:

- `wlac -d wtp [SN|name] [reset]` --> list or reset wtp info(data)
- `wlac -d vap [bssid] [reset]` --> list or reset vap info(data)
- `wlac -d sta [mac] [reset]` --> list or reset sta info(data)

Reference

This chapter provides some reference information pertaining to wireless networks.

FortiAP web-based manager

You can access the FortiAP unit's built-in web-based manager. This is useful to adjust settings that are not available through the FortiGate unit's WiFi Controller. Logging into the FortiAP web-based manager is similar to logging into the FortiGate web-based manager.

System information

Status

The **Status** section provides information about the FortiAP unit.

Host Name	FAP22B3U11005354 [Change]
Serial Number	FAP22B3U11005354
Region Code	A
Firmware Version	FortiAP-220B v5.0,build064,140117 (GA) [Update]
Network Status	0.0.0.0/0.0.0.0/0.0.0.0 (Mon Jun 2 12:50:05 2014)
System Time	Tue May 27 13:00:39 2014
Current Administrator	admin [Change Password]
System Configuration	Last Backup: N/A [Backup] [Restore]
Uptime	7 day(s) 2 hour(s) 4 min(s)
CPU Usage	<div><div></div></div> 1%
Memory Usage	<div><div></div></div> 35%
AC Discovery Status	Discovering AC ...

You can:

- Select **Change** to change the **Host Name**.
- Select **Update** in **Firmware Version** to upload a new FortiAP firmware file from your computer.
- Select **Change Password** to change the administrator password.
- Select **Backup** to save the current FortiAP configuration as a file on your computer.
- Select **Restore** to load a configuration into your FortiAP unit from a file on your computer.

Network configuration

Select DHCP or select Static and specify the IP address, netmask, and gateway IP address. **Administrative Access** settings affect access after the FortiAP has been authorized. By default, **HTTP** access needed to access the FortiAP web-based manager is enabled, but **Telnet** access is not enabled.

Connectivity

These settings determine how the FortiAP unit connects to the FortiGate WiFi controller.

Uplink	Ethernet - wired connection to the FortiGate unit (default) Mesh - WiFi mesh connection Ethernet with mesh backup support
Mesh AP SSID	Enter the SSID of the mesh root. Default: fortinet.mesh.root
Mesh AP Password	Enter password for the mesh SSID.
Ethernet Bridge	Bridge the mesh SSID to the FortiAP Ethernet port. This is available only when Uplink is Mesh .

WTP configuration

AC Discovery Type settings affect how the FortiAP unit discovers a FortiGate WiFi controller. By default, this is set to Auto which causes the FortiAP unit to cycle through all of the discovery methods until successful. For more information see Controller discovery methods.

AC Discovery Type	Static, DHCP, DNS, Broadcast, Multicast, Auto
AC Control Port	Default port is 5246.
AC IP Address 1 AC IP Address 2 AC IP Address 3	You enter up to three WiFi controller IP addresses for static discovery. Routing must be properly configured in both directions.
AC Host Name 1 AC Host Name 2 AC Host Name 3	As an alternative to AC IP addresses, you can enter their fully qualified domain names (FQDNs).
AC Discovery Multicast Address	224.0.1.140
AC Discovery DHCP Option Code	When using DHCP discovery, you can configure the DHCP server to provide the controller address. By default the FortiAP unit expects this in option 138.

AC Data Channel Security by default accepts either DTLS-encrypted or clear text data communication with the WiFi controller. You can change this setting to require encryption or to use clear text only.

Wireless information

The Wireless Information page provides current information about the operation of the radios and the type Uplink in use.

Wireless radio channels

IEEE 802.11a/n channels

The following table lists the channels supported on FortiWiFi products that support the IEEE 802.11a and 802.11n wireless standards. 802.11a is available on FortiWiFi models 60B and higher. 802.11n is available on FortiWiFi models 80CM and higher.

All channels are restricted to indoor usage except in the Americas, where both indoor and outdoor use is permitted on channels 52 through 64 in the United States.

IEEE 802.11a/n (5-GHz Band) channel numbers

Channel number	Frequency (MHz)	Regulatory Areas				
		Americas	Europe	Taiwan	Singapore	Japan
34	5170					•
36	5180	•	•		•	
38	5190					
40	5200	•	•		•	•
42	5210					
44	5220	•	•		•	•
46	5230					
48	5240	•	•		•	•
149	5745	•		•	•	
153	5765	•		•	•	
157	5785	•		•	•	
161	5805	•		•	•	
165	5825	•			•	

IEEE 802.11b/g/n channel numbers

The following table lists IEEE 802.11b/g/n channels. All FortiWiFi units support 802.11b and 802.11g. Newer models also support 802.11n.

Mexico is included in the Americas regulatory domain. Channels 1 through 8 are for indoor use only. Channels 9 through 11 can be used indoors and outdoors. You must make sure that the channel number complies with the regulatory standards of Mexico.

IEEE 802.11b/g/n (2.4-GHz Band) channel numbers

Channel number	Frequency (MHz)	Regulatory Areas			
		Americas	EMEA	Israel	Japan
1	2412	•	•	indoor	•
2	2417	•	•	indoor	•
3	2422	•	•	indoor	•
4	2427	•	•	indoor	•
5	2432	•	•	•	•
6	2437	•	•	•	•
7	2442	•	•	•	•
8	2447	•	•	•	•
9	2452	•	•	•	•
10	2457	•	•	•	•
11	2462	•	•	•	•
12	2467		•	•	•
13	2472		•	•	•
14	2484				b only

View all country and regcodes/regulatory domains

The following CLI command can be entered to view a list of the country and regcodes/regulatory Domains supported by Fortinet:

```
cw_diag -c all-countries
```

Below is a table showing a sample of the list displayed by entering this command:

Country-code	Region-code	Domain	ISO-name	Name
0	A	FCC3 & FCCA	NA	NO_COUNTRY_SET

Country-code	Region-code	Domain	ISO-name	Name
8	W	NULL1 & WORLD	AL	ALBANIA
12	W	NULL1 & WORLD	DZ	ALGERIA
16	A	FCC3 & FCCA	AS	AMERICAN SAMOA
...

WiFi event types

Event type	Description
rogue-ap-detected	A rogue AP has been detected (generic).
rogue-ap-off-air	A rogue AP is no longer detected on the RF side.
rogue-ap-on-wire	A rogue AP has been detected on wire side (connected to AP or controller L2 network).
rogue-ap-off-wire	A rogue AP is no longer detected on wire.
rogue-ap-on-air	A rogue AP has been detected on the RF side.
fake-ap-detected	A rogue AP broadcasting on the same SSIDs that you have in your managed APs has been detected.
fake-ap-on-air	The above fake AP was detected on the RF side.

FortiAP CLI

The FortiAP CLI controls radio and network operation through the use of variables manipulated with the `cfg` command. There are also diagnostic commands.

The `cfg` command include the following

<code>cfg -s</code>	List variables.
<code>cfg -a var=value</code>	Add or change a variable value.
<code>cfg -c</code>	Commit the change to flash.
<code>cfg -x</code>	Reset settings to factory defaults.
<code>cfg -r var</code>	Remove variable.

<code>cfg -e</code>	Export variables.
<code>cfg -h</code>	Display help for all commands.

The configuration variables are:

Var	Description and Values
AC_CTL_PORT	WiFi Controller control (CAPWAP) port. Default 5246.
AC_DATA_CHAN_SEC	Data channel security. 0 - Clear text 1 - DTLS (encrypted) 2 - Accept either DTLS or clear text (default)
AC_DISCOVERY_TYPE	1 - Static. Specify WiFi Controllers 2 - DHCP 3 - DNS 5 - Broadcast 6 - Multicast 0 - Cycle through all of the discovery types until successful.
AP_IPADDR AP_NETMASK IPGW	These variables set the FortiAP unit IP address, netmask and default gateway when ADDR_MODE is STATIC. Default 192.168.1.2 255.255.255.0, gateway 192.168.1.1.
AC_HOSTNAME_1 AC_HOSTNAME_2 AC_HOSTNAME_3	WiFi Controller host names for static discovery.
AC_IPADDR_1 AC_IPADDR_2 AC_IPADDR_3	WiFi Controller IP addresses for static discovery.
AC_DISCOVERY_DHCP_OPTION_CODE	Option code for DHCP server. Default 138.
AC_DISCOVERY_MC_ADDR	Multicast address for controller discovery. Default 224.0.1.140.

Var	Description and Values
ADDR_MODE	How the FortiAP unit obtains its IP address and netmask. DHCP - FortiGate interface assigns address. STATIC - Specify in AP_IPADDR and AP_NETMASK. Default is DHCP.
ADMIN_TIMEOUT	Administrative timeout in minutes. Applies to Telnet and web-based manager sessions. Default is 5 minutes.
AP_MGMT_VLAN_ID	Non-zero value applies VLAN ID for unit management. Default: 0.
AP_MODE	FortiAP operating mode. 0 - Thin AP (default) 2 - Unmanaged Site Survey mode. See SURVEY variables.
BAUD_RATE	Console data rate: 9600, 19200, 38400, 57600, or 115200 baud.
DNS_SERVER	DNS Server for clients. If ADDR_MODE is DHCP the DNS server is automatically assigned.
FIRMWARE_UPGRADE	Default is 0.
HTTP_ALLOW	Access to FortiAP web-based manager 1 - Yes (default), 0 - No.
LED_STATE	Enable/disable status LEDs. 0 - LEDs enabled, 1 - LEDs disabled, 2 - follow AC setting.
LOGIN_PASSWD	Administrator login password. By default this is empty.
STP_MODE	Spanning Tree Protocol. 0 is off. 1 is on.
TELNET_ALLOW	By default (value 0), Telnet access is closed when the FortiAP unit is authorized. Set value to 1 to keep Telnet always available.
WTP_LOCATION	Optional string describing AP location.
Mesh variables	

Var	Description and Values
MESH_AP_BGSCAN	<p>Enable or disable background mesh root AP scan.</p> <p>0 - Disabled</p> <p>1 - Enabled</p>
MESH_AP_BGSCAN_RSSI	<p>If the root AP's signal is weak, and lower than the received signal strength indicator (RSSI) threshold, the WiFi driver will immediately start a new round scan and ignore the configured MESH_AP_BGSCAN_PERIOD delays. Set the value between 0-127.</p> <p>After the new round scan is finished, a scan done event is passed to wtp daemon to trigger roaming.</p>
MESH_AP_BGSCAN_PERIOD	Time in seconds that a delay period occurs between scans. Set the value between 1-3600.
MESH_AP_BGSCAN_IDLE	Time in milliseconds. Set the value between 0-1000.
MESH_AP_BGSCAN_INTV	Time in milliseconds between channel scans. Set the value between 200-16000.
MESH_AP_BGSCAN_DUR	Time in milliseconds that the radio will continue scanning the channel. Set the value between 10-200.
MESH_AP_SCANCHANLIST	Specify those channels to be scanned.
MESH_AP_TYPE	<p>Type of communication for backhaul to controller:</p> <p>0 - Ethernet (default)</p> <p>1 - WiFi mesh</p> <p>2 - Ethernet with mesh backup support</p>
MESH_AP_SSID	SSID for mesh backhaul. Default: fortinet.mesh.root
MESH_AP_BSSID	WiFi MAC address
MESH_AP_PASSWD	Pre-shared key for mesh backhaul.
MESH_ETH_BRIDGE	<p>1 - Bridge mesh WiFi SSID to FortiAP Ethernet port. This can be used for point-to-point bridge configuration. This is available only when MESH_AP_TYPE =1.</p> <p>0 - No WiFi-Ethernet bridge (default).</p>

Var	Description and Values
MESH_MAX_HOPS	Maximum number of times packets can be passed from node to node on the mesh. Default is 4.
The following factors are summed and the FortiAP associates with the lowest scoring mesh AP.	
MESH_SCORE_HOP_WEIGHT	Multiplier for number of mesh hops from root. Default 50.
MESH_SCORE_CHAN_WEIGHT	AP total RSSI multiplier. Default 1.
MESH_SCORE_RATE_WEIGHT	Beacon data rate multiplier. Default 1.
MESH_SCORE_BAND_WEIGHT	Band weight (0 for 2.4GHz, 1 for 5GHz) multiplier. Default 100.
MESH_SCORE_RSSI_WEIGHT	AP channel RSSI multiplier. Default 100.
Survey variables	
SURVEY_SSID	SSID to broadcast in site survey mode (AP_MODE=2).
SURVEY_TX_POWER	Transmitter power in site survey mode (AP_MODE=2).
SURVEY_CH_24	Site survey transmit channel for the 2.4Ghz band (default 6).
SURVEY_CH_50	Site survey transmit channel for the 5Ghz band (default 36).
SURVEY_BEACON_INTV	Site survey beacon interval. Default 100msec.



Previously, FortiAP accepted Telnet and HTTP connection to any virtual interfaces that have an IP address. For security reasons, Telnet and HTTP access are now limited to br0 or br.vlan for AP_MGMT_VLAN_ID.

Diagnose commands include:

<code>cw_diag help</code>	Display help for all diagnose commands.
<code>cw_diag uptime</code>	Show daemon uptime.
<code>cw_diag --tlog <on off></code>	Turn on/off telnet log message.
<code>cw_diag --clog <on off></code>	Turn on/off console log message.
<code>cw_diag baudrate [9600 19200 38400 57600 115200]</code>	Set the console baud rate.

<code>cw_diag plain-ctl [0 1]</code>	Show or change current plain control setting.
<code>cw_diag sniff-cfg ip port</code>	Set sniff server ip and port.
<code>cw_diag sniff [0 1 2]</code>	Enable/disable sniff packet.
<code>cw_diag stats wl_intf</code>	Show wl_intf status.
<code>cw_diag admin-timeout [30]</code>	Set shell idle timeout in minutes.
<code>cw_diag -c wtp-cfg</code>	Show current wtp config parameters in control plane.
<code>cw_diag -c radio-cfg</code>	Show current radio config parameters in control plane.
<code>cw_diag -c vap-cfg</code>	Show current vaps in control plane.
<code>cw_diag -c ap-rogue</code>	Show rogue APs pushed by AC for on-wire scan.
<code>cw_diag -c sta-rogue</code>	Show rogue STAs pushed by AC for on-wire scan.
<code>cw_diag -c arp-req</code>	Show scanned arp requests.
<code>cw_diag -c ap-scan</code>	Show scanned APs.
<code>cw_diag -c sta-scan</code>	Show scanned STAs.
<code>cw_diag -c sta-cap</code>	Show scanned STA capabilities.
<code>cw_diag -c wids</code>	Show scanned WIDS detections.
<code>cw_diag -c darrp</code>	Show darrp radio channel.
<code>cw_diag -c mesh</code>	Show mesh status.
<code>cw_diag -c mesh-veth-acinfo</code>	Show mesh veth ac info, and mesh ether type.
<code>cw_diag -c mesh-veth-vap</code>	Show mesh veth vap.
<code>cw_diag -c mesh-veth-host</code>	Show mesh veth host.
<code>cw_diag -c mesh-ap</code>	Show mesh ap candidates.
<code>cw_diag -c scan-clr-all</code>	Flush all scanned AP/STA/ARPs.
<code>cw_diag -c ap-suppress</code>	Show suppressed APs.
<code>cw_diag -c sta-deauth</code>	De-authenticate an STA.



Link aggregation can also be set in the CLI. Link aggregation is used to combine multiple network connections in parallel in order to increase throughput beyond what a single connection could sustain.

- FortiAP 320B and 320C models are supported.
 - FortiAP 112B and 112D models **cannot** support link aggregation.
 - NPI FAP-S3xxCR and "wave2" FAP/FAP-S models will have link aggregation feature via synchronization with regular FortiAP trunk build.
-

Chapter 10 - Hardening your FortiGate

This guide describes some of the techniques used to harden (improve the security of) FortiGate devices and FortiOS.

This document contains the following sections:

- [Building security into FortiOS](#)
- [FortiOS ports and protocols](#)
- [Security best practices](#)

Building security into FortiOS

The FortiOS operating system, FortiGate hardware devices, and FortiOS virtual machines (VMs) are built with security in mind, so many security features are built into the hardware and software. Fortinet maintains an ISO:9001 certified software and hardware development processes to ensure that FortiOS and FortiGate products are developed in a secure manner.

Boot PROM and BIOS security

The boot PROM and BIOS in FortiGate hardware devices use Fortinet's own FortiBootLoader that is designed and controlled by Fortinet. FortiBootLoader is a secure, proprietary BIOS for all FortiGate appliances. FortiGate physical devices always boot from FortiBootLoader.

FortiOS kernel and user processes

FortiOS is a multi-process operating system with kernel and user processes. The FortiOS kernel runs in a privileged hardware mode while higher-level applications run in user mode. FortiOS is a closed system that does not allow the loading or execution of third-party code in the FortiOS user space. All non-essential services, packages, and applications are removed.

FortiGate appliances with SD drives are encrypted to prevent unauthorized access to data.

Administration access security

This section describes FortiOS and FortiGate administration access security features.

Admin administrator account

All FortiGate firewalls ship with a default administrator account called admin. By default, this account does not have a password. FortiOS allows administrators to add a password for this account or to remove the account and create new custom super_admin administrator accounts.

For more information, see [Rename the admin administrator account on page 1234](#).

Secure password storage

User and administrator passwords are stored securely on the system in an encrypted format. The encryption hash used for admin account passwords is SHA256/SHA1. The value that is seen in the configuration file is the Base64 encoded hash value. For example:

```
config system admin
  edit "admin"
    set accprofile "super_admin"
    set vdom "root"
    set password ENC SH2nlSm9QL9tapcHPXIqAXvX7vBJuuqu22hpa0JX0sBuKIo7z2g0Kz/+0KyH4E=
  next
end
```

Pre-shared keys in IPSec phase-1 configurations are stored in plain text. In the configuration file these pre-shared keys are encoded. The encoding consists of encrypting the password with a fixed key using DES (AES in FIPS mode) and then Base64 encoding the result.

Maintainer account

Administrators with physical access to a FortiGate appliance can use a console cable and a special administrator account called maintainer to log into the CLI. When enabled, the maintainer account can be used to log in from the console after a hard reboot. The password for the maintainer account is bcpb followed by the FortiGate serial number. An administrator has 60-seconds to complete this login. See [Resetting a lost Admin password](#) on the Fortinet Cookbook for details.

The only action the maintainer account has permissions to perform is to reset the passwords of super_admin accounts. Logging in with the maintainer account requires rebooting the FortiGate. FortiOS generates event log messages when you login with the maintainer account and for each password reset.

The maintainer account is enabled by default; however, there is an option to disable this feature. The maintainer account can be disabled using the following command:

```
config system global
  set admin-maintainer disable
end
```



If you disable this feature and lose your administrator passwords you will no longer be able to log into your FortiGate.

Administrative access security

Secure administrative access features:

- SSH, Telnet, and SNMP are disabled by default. If required, these admin services must be explicitly enabled on each interface from the GUI or CLI.
- SSHv1 is disabled by default. SSHv2 is the default version.
- SSLv3 and TLS1.0 are disabled by default. TLSv1.1 and TLSv1.2 are the SSL versions enabled by default for HTTPS admin access.
- HTTP is disabled by default, except on dedicated MGMT, DMZ, and predefined LAN interfaces. HTTP redirect to HTTPS is enabled by default.
- The `strong-crypto` global setting is enabled by default and configures FortiOS to use strong ciphers (AES, 3DES) and digest (SHA1) for HTTPS/SSH/TLS/SSL functions.
- SCP is disabled by default. Enabling SCP allows downloading the configuration file from the FortiGate as an alternative method of backing up the configuration file. To enable SCP:

```
config system global
  set admin-scp enable
end
```

- DHCP is enabled by default on the dedicated MGMT interface and on the predefined LAN port (defined on some FortiGate models).
- The default management access configuration for FortiGate models with dedicated MGMT, DMZ, WAN, and LAN interfaces is shown below. Outside of the interfaces listed below, management access must be explicitly enabled on interfaces – management services are enabled on specific interfaces and not globally.
 - Dedicated management interface
 - Ping
 - FMG-Access (fgfm)
 - CAPWAP

- HTTPS
- HTTP
- Dedicated WAN1/WAN2 interface
 - Ping
 - FMG-Access (fgfm)
- Dedicated DMZ interface
 - Ping
 - FMG-Access (fgfm)
 - CAPWAP
 - HTTPS
 - HTTP
- Dedicated LAN interface
 - Ping
 - FMG-Access (fgfm)
 - CAPWAP
 - HTTPS
 - HTTP

Network security

This section describes FortiOS and FortiGate network security features.

Network interfaces

The following are disabled by default on each FortiGate interface:

- Broadcast forwarding
- STP forwarding
- VLAN forwarding
- L2 forwarding
- Netbios forwarding
- Ident accept

For more information, see [Disable unused protocols on interfaces on page 1237](#).

TCP sequence checking

FortiOS uses TCP sequence checking to ensure a packet is part of a TCP session. By default, anti-replay protection is strict, which means that if a packet is received with sequence numbers that fall out of the expected range, FortiOS drops the packet. Strict anti-replay checking performs packet sequence checking and ICMP anti-replay checking with the following criteria:

- The SYN, FIN, and RST bit cannot appear in the same packet.
- FortiOS does not allow more than 1 ICMP error packet to go through before it receives a normal TCP or UDP packet.
- If FortiOS receives an RST packet, FortiOS checks to determine if its sequence number in the RST is within the un-ACKed data and drops the packet if the sequence number is incorrect.
- For each new session, FortiOS checks to determine if the TCP sequence number in a SYN packet has been calculated correctly and started from the correct value.

Reverse path forwarding

FortiOS implements a mechanism called Reverse Path Forwarding (RPF), or Anti Spoofing, to block an IP packet from being forwarded if its source IP does not:

- belong to a locally attached subnet (local interface), or
- be in the routing domain of the FortiGate from another source (static route, RIP, OSPF, BGP).

If those conditions are not met, FortiOS silently drops the packet.

FIPS and Common Criteria

FortiOS has received NDPP, EAL2+, and EAL4+ based FIPS and Common Criteria certifications. Common Criteria evaluations involve formal rigorous analysis and testing to examine security aspects of a product or system. Extensive testing activities involve a comprehensive and formally repeatable process, confirming that the security product functions as claimed by the manufacturer. Security weaknesses and potential vulnerabilities are specifically examined during an evaluation.

To see Fortinet's complete history of FIPS/CC certifications go to the following URL and add Fortinet to the Vendor field:

<https://csrc.nist.gov/projects/cryptographic-module-validation-program/validated-modules/search>

PSIRT advisories

The FortiGuard Labs Product Security Incident Response Team (PSIRT) continually tests and gathers information about Fortinet hardware and software products, looking for vulnerabilities and weaknesses. Any such findings are fed back to Fortinet's development teams and serious issues are described along with protective solutions. The PSIRT regulatory releases PSIRT advisories when issues are found and corrected. Advisories are listed at <https://www.fortiguard.com/psirt>.

FortiOS ports and protocols

Communication to and from FortiOS is strictly controlled and only selected ports are opened for supported functionality such as administrator logins and communication with other Fortinet products or services.

Accessing FortiOS using an open port is protected by authentication, identification, and encryption requirements. As well, ports are only open if the feature using them is enabled.

FortiOS open ports

The following diagram and tables shows the incoming and outgoing ports that are potentially opened by FortiOS. For more details about open ports and the communication protocols that FortiOS uses, see the document [Fortinet Communication Ports and Protocols](#).

(missing or bad snippet)

Closing open ports

You can close open ports by disabling the feature that opens them. For example, if FortiOS is not managing a FortiAP then the CAPWAP feature for managing FortiAPs can be disabled, closing the CAPWAP port.

The following sections of this document described a number of options for closing open ports:

- [Use local-in policies to close open ports or restrict access on page 1238](#)
- [Disable unused protocols on interfaces on page 1237](#)

Security best practices

This chapter describes some techniques and best practices that you can use to improve FortiOS security.

Install the FortiGate unit in a physically secure location

A good place to start with is physical security. Install your FortiGate in a secure location, such as a locked room or one with restricted access. A restricted location prevents unauthorized users from getting physical access to the device.

If unauthorized users have physical access, they can disrupt your entire network by disconnecting your FortiGate (either by accident or on purpose). They could also connect a console cable and attempt to log into the CLI. Also, when a FortiGate unit reboots, a person with physical access can interrupt the boot process and install different firmware.

Register your product with Fortinet Support

You need to register your Fortinet product with Fortinet Support to receive customer services, such as firmware updates and customer support. You must also register your product for FortiGuard services, such as up-to-date antivirus and IPS signatures. To register your product the [Fortinet Support](#) website.

Keep your FortiOS firmware up to date

Always keep FortiOS up to date. The most recent version is the most stable and has the most bugs fixed and vulnerabilities removed. Fortinet periodically updates the FortiGate firmware to include new features and resolve important issues.

After you register your FortiGate, you can receive notifications on FortiGate GUI about firmware updates. You can update the firmware directly from the GUI or by downloading firmware updates from the [Fortinet Support](#) website.

Before you install any new firmware, be sure to follow these steps:

- Review the release notes for the latest firmware release.
- Review the [Supported Upgrade Paths](#) guide to determine the best path to take from your current version of FortiOS to the latest version.
- Back up the current configuration.

Only FortiGate administrators who have read and write privileges can upgrade the FortiOS firmware.

System administrator best practices

This section describes a collection of changes you can implement to make administrative access to the GUI and CLI more secure.

Disable administrative access to the external (Internet-facing) interface

When possible, don't allow administration access on the external (Internet-facing) interface.

To disable administrative access, go to **Network > Interfaces**, edit the external interface and disable HTTPS, PING, HTTP, SSH, and TELNET under **Administrative Access**.

From the CLI:

```
config system interface
  edit <external-interface-name>
    unset allowaccess
  end
```

Allow only HTTPS access to the GUI and SSH access to the CLI

For greater security never allow HTTP or Telnet administrative access to a FortiGate interface, only allow HTTPS and SSH access. You can change these settings for individual interfaces by going to **Network > Interfaces** and adjusting the administrative access to each interface.

From the CLI:

```
config system interface
  edit <interface-name>
    set allowaccess https ssh
  end
```

Require TLS 1.2 for HTTPS administrator access

Use the following command to require TLS 1.2 for HTTPS administrator access to the GUI:

```
config system global
  set admin-https-ssl-versions tlsv1-2
end
```

TLS 1.2 is currently the most secure SSL/TLS supported version for SSL-encrypted administrator access.

Re-direct HTTP GUI logins to HTTPS

Go to **System > Settings > Administrator Settings** and enable **Redirect to HTTPS** to make sure that all attempted HTTP login connections are redirected to HTTPS.

From the CLI:

```
config system global
  set admin-https-redirect enable
end
```

Change the HTTPS and SSH admin access ports to non-standard ports

Go to **System > Settings > Administrator Settings** and change the HTTPS and SSH ports.

You can change the default port configurations for HTTPS and SSH administrative access for added security. To connect to a non-standard port, the new port number must be included in the collection request. For example:

- If you change the HTTPS port to 7734, you would browse to `https://<ip-address>:7734`.
- If you change the SSH port to 2345, you would connect to `ssh admin@<ip-address>:2345`

To change the HTTPS and SSH login ports from the CLI:

```
config system global
  set admin-sport 7734
  set admin-ssh-port 2345
end
```

If you change to the HTTPS or SSH port numbers, make sure your changes do not conflict with ports used for other services.

Maintain short login timeouts

Set the idle timeout to a short time to avoid the possibility of an administrator walking away from their management computer and leaving it exposed to unauthorized personnel.

To set the administrator idle timeout, go to **System > Settings** and enter the amount of time for the **Idle timeout**. A best practice is to keep the default time of 5 minutes.

To set the administrator idle timeout from the CLI:

```
config system global
    set admintimeout 5
end
```

You can use the following command to adjust the grace time permitted between making an SSH connection and authenticating. The range can be between 10 and 3600 seconds, the default is 120 seconds (minutes). By shortening this time, you can decrease the chances of someone attempting a brute force attack from being successful. For example, you could set the time to 30 seconds.

```
config system global
    set admin-ssh-grace-time 30
end
```

Restrict logins from trusted hosts

Setting up trusted hosts for an administrator limits the addresses from where they can log into FortiOS. The trusted hosts configuration applies to most forms of administrative access including HTTPS, SSH, and SNMP. When you identify a trusted host for an administrator account, FortiOS accepts that administrator's login only from one of the trusted hosts. A login, even with proper credentials, from a non-trusted host is dropped.



Even if you have configured trusted hosts, if you have enabled ping administrative access on a FortiGate interface, it will respond to ping requests from any IP address.

To identify trusted hosts, go to **System > Administrators**, edit the administrator account, enable **Restrict login to trusted hosts**, and add up to ten trusted host IP addresses.

To add two trusted hosts from the CLI:

```
config system admin
    edit <administrator-name>
        set trustedhost1 172.25.176.23 255.255.255.255
        set trustedhost2 172.25.177.0 255.255.255.0
    end
```

Trusted host IP addresses can identify individual hosts or subnets. Just like firewall policies, FortiOS searches through the list of trusted hosts in order and acts on the first match it finds. When you configure trusted hosts, start by adding specific addresses at the top of the list. Follow with more general IP addresses. You don't have to add addresses to all of the trusted hosts as long as all specific addresses are above all of the 0.0.0.0 0.0.0.0 addresses.

Set up two-factor authentication for administrators

FortiOS supports FortiToken and FortiToken Mobile 2-factor authentication. FortiToken Mobile is available for iOS and Android devices from their respective application stores.

Every registered FortiGate unit includes two trial tokens for free. You can purchase additional tokens from your reseller or from Fortinet.

To assign a token to an administrator, go to **System > Administrators** and select **Enable Two-factor Authentication** for each administrator.

Create multiple administrator accounts

Rather than allowing all administrators to access FortiOS with the same administrator account, you can create accounts for each person or each role that requires administrative access. This configuration allows you to track the activities of each administrator or administrative role.

If you want administrators to have different functions you can add different administrator profiles. Go to **System > Admin Profiles** and select **Create New**.

Modify administrator account lockout duration and threshold values

By default, the FortiGate sets the number of password retries at three, allowing the administrator a maximum of three attempts to log into their account before locking the account for a set amount of time.

Both the number of attempts (`admin-lockout-threshold`) and the wait time before the administrator can try to enter a password again (`admin-lockout-duration`) can be configured within the CLI.

To configure the lockout options:

```
config system global
  set admin-lockout-threshold <failed_attempts>
  set admin-lockout-duration <seconds>
end
```

The default value of `admin-lockout-threshold` is 3 and the range of values is between 1 and 10. The `admin-lockout-duration` is set to 60 seconds by default and the range of values is between 1 and 4294967295 seconds.

Keep in mind that the higher the lockout threshold, the higher the risk that someone may be able to break into the FortiGate.

Example:

To set the `admin-lockout-threshold` to one attempt and the `admin-lockout-duration` to a five minute duration before the administrator can try to log in again, enter the commands:

```
config system global
  set admin-lockout-threshold 1
  set admin-lockout-duration 300
end
```



If the time span between the first failed login attempt and the `admin-lockout-threshold` failed login attempt is less than `admin-lockout-duration`, the lockout will be triggered.

Rename the admin administrator account

You can improve security by renaming the admin account. To do this, create a new administrator account with the `super_admin` admin profile and log in as that administrator. Then go to **System > Administrators** and edit the admin administrator and change the **User Name**. Renaming the admin account makes it more difficult for an attacker to log into FortiOS.

Add administrator disclaimers

FortiOS can display a disclaimer before or after logging into the GUI or CLI (or both). In either case the administrator must read and accept the disclaimer before they can proceed.

Use the following command to display a disclaimer before logging in:

```
config system global
  set pre-login-banner enable
end
```

Use the following command to display a disclaimer after logging in:

```
config system global
  set post-login-banner enable
end
```

You can customize the replacement messages for these disclaimers by going to **System > Replacement Messages**. Select **Extended View** to view and edit the **Administrator** replacement messages.

From the CLI:

```
config system replacemsg admin pre_admin-disclaimer-text
config system replacemsg admin post_admin-disclaimer-text
```

Global commands for stronger and more secure encryption

This section describes some best practices for employing stronger and more secure encryption.

Turn on global strong encryption

Enter the following command to configure FortiOS to use only strong encryption and allow only strong ciphers (AES, 3DES) and digest (SHA1) for HTTPS, SSH, TLS, and SSL functions.

```
config sys global
  set strong-crypto enable
end
```

Disable MD5 and CBC for SSH

In some cases, you may not be able to enable strong encryption. For example, your FortiGate may be communicating with a system that does not support strong encryption. With `strong-crypto` disabled you can use the following options to prevent SSH sessions with the FortiGate from using less secure MD5 and CBC algorithms:

```
config sys global
  set ssh-hmac-md5 disable
  set ssh-cbc-cipher disable
end
```

Disable static keys for TLS

You can use the following command to prevent TLS sessions from using static keys (AES128-SHA, AES256-SHA, AES128-SHA256, AES256-SHA256):

```
config sys global
    set ssl-static-key-ciphers disable
end
```

Require larger values for Diffie-Hellman exchanges

Larger Diffie-Hellman values result in stronger encryption. Use the following command to force Diffie-Hellman exchanges to use 8192 bit values (the highest configurable DH value).

```
config sys global
    set dh-params 8192
end
```

Disable sending malware statistics to FortiGuard

By default FortiOS periodically sends encrypted malware statistics to FortiGuard. The malware statistics record Antivirus, IPS, or Application Control events. This data is used to improved FortiGuard services. The malware statistics that FortiOS sends do not include any personal or sensitive customer data. The information is not shared with any external parties and is used in accordance with Fortinet's [Privacy Policy](#).

To disable sending malware statistics to FortiGuard, enter the following command:

```
config system global
    set fds-statistics disable
end
```

Disable sending Security Rating statistics to FortiGuard

Security Rating is a Fortinet Security Fabric feature that allows customers to audit their Security Fabric and find and fix security problems. As part of the feature, FortiOS sends your security rating to FortiGuard every time a security rating test runs.

You can opt out of submitting Security Rating scores to FortiGuard. If you opt out you won't be able to see how your organization's scores compare with the scores of other organizations. Instead, an absolute score is shown. Use the following command to disable FortiGuard Security Rating result submission:

```
config system global
    set fortiguard-audit-result-submission disable
end
```

Disable auto USB installation

If USB installation is enabled, an attacker with physical access to a FortiGate could load a new configuration or firmware on the FortiGate using the USB port. You can disable USB installation by entering the following from the CLI:

```
config system auto-install
    set auto-install-config disable
    set auto-install-image disable
end
```

Set system time by synchronizing with an NTP server

For accurate time, use an NTP server to set system time. Synchronized time facilitates auditing and consistency between expiry dates used in expiration of certificates and security protocols.

From the GUI go to **System > Settings > System Time** and select **Synchronize with NTP Server**. By default, this causes FortiOS to synchronize with Fortinet's FortiGuard secure NTP server.

From the CLI you can use one or more different NTP servers:

```
config system ntp
  set type custom
  set ntpsync enable
  config ntpserver
    edit 1
      set server <ntp-server-ip>
    next
    edit 2
      set server <other-ntp-server-ip>
  end
```

Disable the maintainer admin account

Administrators with physical access to a FortiGate appliance can use a console cable and a special administrator account called maintainer to log into the CLI without a password. This feature allows you to log into a FortiGate if you have lost all administrator passwords. See [Resetting a lost Admin password](#) on the Fortinet Cookbook for details.

The maintainer account can be disabled using the following command:

```
config system global
  set admin-maintainer disable
end
```



If you disable this feature and lose your administrator passwords you will no longer be able to log into your FortiGate.

Enable password policies

Go to **System > Settings > Password Policy**, to create a password policy that all administrators must follow. Using the available options you can define the required length of the password, what it must contain (numbers, upper and lower case, and so on) and an expiry time.

Use the password policy feature to make sure all administrators use secure passwords that meet your organization's requirements.

Configure auditing and logging

For optimum security go to **Log & Report > Log Settings** enable **Event Logging**. For best results send log messages to FortiAnalyzer or FortiCloud.

From FortiAnalyzer or FortiCloud, you can view reports or system event log messages to look for system events that may indicate potential problems. You can also view system events by going to **FortiView > System Events**.

Establish an auditing schedule to routinely inspect logs for signs of intrusion and probing.

Encrypt logs sent to FortiAnalyzer/FortiManager

To keep information in log messages sent to FortiAnalyzer private, go to **Log & Report > Log Settings** and when you configure Remote Logging to FortiAnalyzer/FortiManager select **Encrypt log transmission**.

From the CLI.

```
config log {fortianalyzer | fortianalyzer2 | fortianalyzer3} setting
  set enc-algorithm high
end
```

Disable unused interfaces

To disable an interface from the GUI, go to **Network > Interfaces**. Edit the interface to be disabled and set **Interface State** to **Disabled**.

From the CLI, to disable the port21 interface:

```
config system interface
  edit port21
    set status down
  end
```

Disable unused protocols on interfaces

You can use the `config system interface` command to disable unused protocols that attackers may attempt to use to gather information about a FortiGate unit. Many of these protocols are disabled by default. Using the `config system interface` command you can see the current configuration of each of these options for the selected interface and then choose to disable them if required.

```
config system interface
  edit <interface-name>
    set dhcp-relay-service disable
    set pptp-client disable
    set arpforward disable
    set broadcast-forward disable
    set l2forward disable
    set icmp-redirect disable
    set vlanforward disable
    set stpforward disable
    set ident-accept disable
    set ipmac disable
    set netbios-forward disable
    set security-mode none
    set device-identification disable
    set lldp-transmission disable
  end
```

Option	Description
dhcp-relay-service	Disable the DHCP relay service.
pptp-client	Disable operating the interface as a PPTP client.
arpforward	Disable ARP forwarding.
broadcast-forward	Disable forwarding broadcast packets.
l2forward	Disable layer 2 forwarding.
icmp-redirect	Disable ICMP redirect.
vlanforward	Disable VLAN forwarding.
stpforward	Disable STP forwarding.
ident-accept	Disable authentication for this interface. The interface will not respond to a connection with an authentication prompt.
ipmac	Disable IP/MAC binding.
netbios-forward	Disable NETBIOS forwarding.
security-mode	Set to <code>none</code> to disable captive portal authentication. The interface will not respond to a connection with a captive portal.
device-identification	Disable device identification.
lldp-transmission	Disable link layer discovery (LLDP).

Use local-in policies to close open ports or restrict access

You can also use local-in policies to close open ports or otherwise restrict access to FortiOS.

Close ICMP ports

Use the following command to close all ICMP ports on the WAN1 interface. The following example blocks traffic that matches the ICMP_ANY firewall service.

```
config firewall local-in-policy
  edit 1
    set intf wan1
    set srcaddr all
    set dstaddr all
    set action deny
    set service ICMP_ANY
    set schedule always
  end
```

Close the BGP port

Use the following command to close the BGP port on the wan1 interface. The following example blocks traffic that matches the BGP firewall service.

```
config firewall local-in-policy
  edit 1
    set intf wan1
    set srcaddr all
    set dstaddr all
    set action deny
    set service BGP
    set schedule always
  end
```

Chapter 11 - Hardware Acceleration

This FortiOS Handbook chapter contains the following sections:

[Hardware acceleration overview](#) describes the capabilities of FortiGate content processors (CPs), security processors (SPs) and network processors (NPs). This chapter also describes how to determine the hardware acceleration components installed in your FortiGate unit and contains some configuration details and examples.

[NP6 and NP6lite acceleration](#) describes the FortiGate NP6 network processor.

[FortiGate NP6 architectures](#) contains details about the network processing architectures of FortiGate units that contain NP6 processors.

[FortiGate NP6lite architectures on page 1326](#) contains details about the network processing architectures of FortiGate units that contain NP6Lite processors.

[NP4 and NP4Lite acceleration](#) describes the FortiGate NP4 network processor.

[FortiGate NP4 architectures](#) contains details about the network processing architectures of FortiGate units that contain NP4 processors.

What's new in FortiOS 6.0.2

The following list contains new Hardware Acceleration features added in FortiOS 6.0.2. Click on a link to navigate to that section for further information.

- Per-session accounting for NP6Lite processors, see [Enabling per-session accounting for offloaded NP6 and NP6lite sessions on page 1271](#).

What's new in FortiOS 6.0

The following list contains new Hardware Acceleration features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- New options for optimizing FortiGate-3960E and 3980E IPsec VPN performance, see [Optimizing FortiGate-3960E and 3980E IPsec VPN performance on page 1275](#).

Hardware acceleration overview

Most FortiGate models have specialized acceleration hardware, (called Security Processing Units (SPUs)) that can offload resource intensive processing from main processing (CPU) resources. Most FortiGate units include specialized content processors (CPs) that accelerate a wide range of important security processes such as virus scanning, attack detection, encryption and decryption. (Only selected entry-level FortiGate models do not include a CP processor.) Many FortiGate models also contain security processors (SPs) that accelerate processing for specific security features such as IPS and network processors (NPs) that offload processing of high volume network traffic.

Content processors (CP4, CP5, CP6, CP8, and CP9)

Most FortiGate models contain Security Processing Unit (SPU) Content Processors (CPs) that accelerate many common resource intensive security related processes. CPs work at the system level with tasks being offloaded to them as determined by the main CPU. Capabilities of the CPs vary by model. Newer FortiGate units include CP8 and CP9 processors. Older CP versions still in use in currently operating FortiGate models include the CP4, CP5, and CP6.

CP9 capabilities

The CP9 content processor provides the following services:

- Flow-based inspection (IPS, application control etc.) pattern matching acceleration with over 10Gbps throughput
 - IPS pre-scan
 - IPS signature correlation
 - Full match processors
- High performance VPN bulk data engine
 - IPsec and SSL/TLS protocol processor
 - DES/3DES/AES128/192/256 in accordance with FIPS46-3/FIPS81/FIPS197
 - MD5/SHA-1/SHA256/384/512-96/128/192/256 with RFC1321 and FIPS180
 - HMAC in accordance with RFC2104/2403/2404 and FIPS198
 - ESN mode
 - GCM support for NSA "Suite B" (RFC6379/RFC6460) including GCM-128/256; GMAC-128/256
- Key Exchange Processor that supports high performance IKE and RSA computation
 - Public key exponentiation engine with hardware CRT support
 - Primary checking for RSA key generation
 - Handshake accelerator with automatic key material generation
 - True Random Number generator
 - Elliptic Curve support for NSA "Suite B"
 - Sub public key engine (PKCE) to support up to 4096 bit operation directly (4k for DH and 8k for RSA with CRT)
- DLP fingerprint support
 - TTTD (Two-Thresholds-Two-Divisors) content chunking
 - Two thresholds and two divisors are configurable

CP8 capabilities

The CP8 content processor provides the following services:

- Flow-based inspection (IPS, application control etc.) pattern matching acceleration
- High performance VPN bulk data engine
 - IPsec and SSL/TLS protocol processor
 - DES/3DES/AES in accordance with FIPS46-3/FIPS81/FIPS197
 - ARC4 in compliance with RC4
 - MD5/SHA-1/SHA256 with RFC1321 and FIPS180
 - HMAC in accordance with RFC2104/2403/2404 and FIPS198
- Key Exchange Processor support high performance IKE and RSA computation
 - Public key exponentiation engine with hardware CRT support
 - Primarily checking for RSA key generation
 - Handshake accelerator with automatic key material generation
 - Random Number generator compliance with ANSI X9.31
 - Sub public key engine (PKCE) to support up to 4096 bit operation directly
- Message authentication module offers high performance cryptographic engine for calculating SHA256/SHA1/MD5 of data up to 4G bytes (used by many applications)
- PCI express Gen 2 four lanes interface
- Cascade Interface for chip expansion

CP6 capabilities

- Dual content processors
- FIPS-compliant DES/3DES/AES encryption and decryption
- SHA-1 and MD5 HMAC with RFC1321 and FIPS180
- HMAC in accordance with RFC2104/2403/2404 and FIPS198
- IPsec protocol processor
- High performance IPsec engine
- Random Number generator compliance with ANSI X9.31
- Key exchange processor for high performance IKE and RSA computation
- Script Processor
- SSL/TLS protocol processor for SSL content scanning and SSL acceleration

CP5 capabilities

- FIPS-compliant DES/3DES/AES encryption and decryption
- SHA-1 and MD5 HMAC with RFC1321/2104/2403/2404 and FIPS180/FIPS198
- IPsec protocol processor
- High performance IPsec Engine
- Random Number generator compliant with ANSI X9.31
- Public Key Crypto Engine supports high performance IKE and RSA computation
- Script Processor

CP4 capabilities

- FIPS-compliant DES/3DES/AES encryption and decryption
- SHA-1 and MD5 HMAC
- IPsec protocol processor
- Random Number generator
- Public Key Crypto Engine
- Content processing engine
- ANSI X9.31 and PKCS#1 certificate support

Determining the content processor in your FortiGate unit

Use the `get hardware status` CLI command to determine which content processor your FortiGate unit contains. The output looks like this:

```
get hardware status
Model name: FortiGate-100D
ASIC version: CP8
ASIC SRAM: 64M
CPU: Intel(R) Atom(TM) CPU D525 @ 1.80GHz
Number of CPUs: 4
RAM: 1977 MB
Compact Flash: 15331 MB /dev/sda
Hard disk: 15272 MB /dev/sda
USB Flash: not available
Network Card chipset: Intel(R) PRO/1000 Network Connection (rev.0000)
Network Card chipset: bcm-sw Ethernet driver 1.0 (rev.)
```

The ASIC version line lists the content processor model number.

Viewing SSL acceleration status

You can view the status of SSL acceleration using the following command:

```
get vpn status ssl hw-acceleration-status
Acceleration hardware detected: kxp=on cipher=on
```

Where kxp means key exchange acceleration.

Disabling CP offloading for firewall policies

If you want to completely disable offloading to CP processors for test purposes or other reasons, you can do so in security policies. Here are some examples:

For IPv4 security policies.

```
config firewall policy
edit 1
set auto-asic-offload disable
end
```

For IPv6 security policies.

```
config firewall policy6
edit 1
set auto-asic-offload disable
end
```

For multicast security policies.

```
config firewall multicast-policy
edit 1
set auto-asic-offload disable
end
```



Disabling `auto-asic-offload` also disables NP offloading.

Security processors (SPs)

FortiGate Security Processing (SP) modules, such as the SP3 but also including the XLP, XG2, XE2, FE8, and CE4, work at both the interface and system level to increase overall system performance by accelerating specialized security processing. You can configure the SP to favor IPS over firewall processing in hostile high-traffic environments.

SP processors include their own IPS engine which is similar to the FortiOS IPS engine but with the following limitations:

- The SP IPS engine does not support SSL deep inspection. When you have SSL deep inspection enabled for a security policy that includes flow-based inspection or IPS, offloading to the SP is disabled and traffic is processed by the FortiGate CPU and CP processors.
- The SP IPS engine does not support FortiGuard Web Filtering. When you enable flow-based FortiGuard Web Filtering on a FortiGate unit with an SP processor, the SP processor cannot perform FortiGuard lookups and web pages fail to load.

The following security processors are available:

- The SP3 (XLP) is built into the FortiGate-5101B and provides IPS acceleration. No special configuration is required. All IPS processing, including traffic accepted by IPv4 and IPv6 traffic policies and IPv4 and IPv6 DoS policies is accelerated by the built-in SP3 processors.
- The FMC-XG2 is an FMC module with two 10Gb/s SPF+ interfaces that can be used on FortiGate-3950B and FortiGate-3951B units.
- The FortiGate-3140B also contains a built-in XG2 using ports 19 and 20.
- The ADM-XE2 is a dual-width AMC module with two 10Gb/s interfaces that can be used on FortiGate-3810A and FortiGate-5001A-DW systems.
- The ADM-FE8 is a dual-width AMC module with eight 1Gb/s interfaces that can be used with the FortiGate-3810A.
- The ASM-CE4 is a single-width AMC module with four 10/100/1000 Mb/s interfaces that can be used on FortiGate-3016B and FortiGate-3810A units.



Traffic is blocked if you enable IPS for traffic passing over inter-VDOM links if that traffic is being offloaded by an SP processor. If you disable SP offloading, traffic will be allowed to flow. You can disable offloading in individual firewall policies by disabling `auto-asic-offload` for those policies. You can also use the following command to disable all IPS offloading:

```
config ips global
set np-accel-mode none
set cp-accel-mode none
end
```

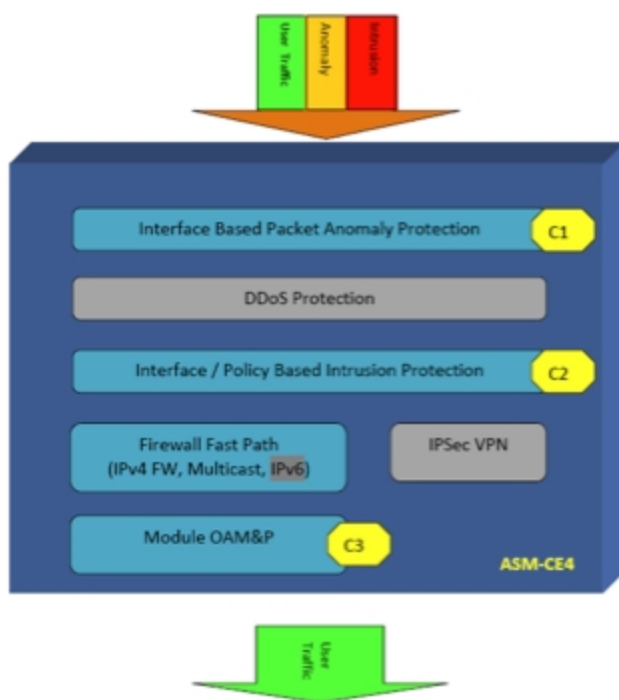
SP processing flow

SP processors provide an integrated high performance fast path multilayer solution for both intrusion protection and firewall functions. The multilayered protection starts from anomaly checking at packet level to ensure each packet is sound and reasonable. Immediately after that, a sophisticated set of interface based packet anomaly protection, DDoS protection, policy based intrusion protection, firewall fast path, and behavior based methods are employed to prevent DDoS attacks from the rest of system.

Then the packets enter an interface/policy based intrusion protection system,

where each packet is evaluated against a set of signatures. The end result is streams of user packets that are free of anomaly and attacks, entering the fast path system for unicast or multicast fast path forwarding.

SP processing flow



Displaying information about security processing modules

You can display information about installed SP modules using the CLI command

```
diagnose npu spm
```

For example, for the FortiGate-5101C:

```
FG-5101C # diagnose npu spm list
Available SP Modules:
```

ID	Model	Slot	Interface
0	xh0	built-in	port1, port2, port3, port4, base1, base2, fabric1, fabric2 eth10, eth11, eth12, eth13 eth14, eth15, eth16, eth17 eth18, eth19

You can also use this command to get more info about SP processing. This example shows how to display details about how the module is processing sessions using the syn proxy.

```
diagnose npu spm dos synproxy <sp_id>
```

This is a partial output of the command:

```
Number of proxied TCP connections : 0
Number of working proxied TCP connections : 0
Number of retired TCP connections : 0
Number of valid TCP connections : 0
Number of attacks, no ACK from client : 0
Number of no SYN-ACK from server : 0
Number of reset by server (service not supported): 0
Number of established session timeout : 0
Client timeout setting : 3 Seconds
Server timeout setting : 3 Seconds
```

Network processors (NP1, NP2, NP3, NP4, NP4Lite, NP6 and NP6Lite)

FortiASIC network processors work at the interface level to accelerate traffic by offloading traffic from the main CPU. Current models contain NP4, NP4Lite, NP6, and NP6lite network processors. Older FortiGate models include NP1 network processors (also known as FortiAccel, or FA2) and NP2 network processors.

The traffic that can be offloaded, maximum throughput, and number of network interfaces supported by each varies by processor model:

- NP6 supports offloading of most IPv4 and IPv6 traffic, IPsec VPN encryption, CAPWAP traffic, and multicast traffic. The NP6 has a maximum throughput of 40 Gbps using 4 x 10 Gbps XAUI or Quad Serial Gigabit Media Independent Interface (QSGMII) interfaces or 3 x 10 Gbps and 16 x 1 Gbps XAUI or QSGMII interfaces. For details about the NP6 processor, see [NP6 and NP6lite acceleration on page 1255](#) and for information about FortiGate models with NP6 processors, see [FortiGate NP6 architectures on page 1276](#).
- NP6lite is similar to the NP6 but with a lower throughput and some functional limitations (for example, the NP6lite does not offload CAPWAP traffic). The NP6lite has a maximum throughput of 10 Gbps using 2x QSGMII and 2x Reduced gigabit media-independent interface (RGMII) interfaces. For details about the NP6 processor, see [NP6Lite processors on page 1257](#) and for information about FortiGate models with NP6 processors, see [FortiGate NP6lite architectures on page 1326](#).
- NP4 supports offloading of most IPv4 firewall traffic and IPsec VPN encryption. The NP4 has a capacity of 20 Gbps through 2 x 10 Gbps interfaces. For details about NP4 processors, see [NP4 and NP4Lite acceleration on page 1328](#) and for information about FortiGate models with NP4 processors, see [FortiGate NP4 architectures on page 1339](#).
- NP4lite is similar to the NP4 but with a lower throughput (but with about half the performance) and some functional limitations.
- NP2 supports IPv4 firewall and IPsec VPN acceleration. The NP2 has a capacity of 2 Gbps through 2 x 10 Gbps interfaces or 4 x 1 Gbps interfaces.
- NP1 supports IPv4 firewall and IPsec VPN acceleration with 2 Gbps capacity. The NP1 has a capacity of 2 Gbps through 2 x 1 Gbps interfaces.
 - The NP1 does not support frames greater than 1500 bytes. If your network uses jumbo frames, you may need to adjust the MTU (Maximum Transmission Unit) of devices connected to NP1 ports. Maximum frame size for NP2, NP4, and NP6 processors is 9216 bytes.
 - For both NP1 and NP2 network processors, ports attached to a network processor cannot be used for firmware installation by TFTP.



Sessions that require proxy-based security features (for example, virus scanning, IPS, application control and so on) are not fast pathed and must be processed by the CPU. Sessions that require flow-based security features can be offloaded to NP4 or NP6 network processors if the FortiGate supports NTurbo.

Determining the network processors installed in your FortiGate

Use either of the following command to list the NP6 processors in your FortiGate unit:

```
get hardware npu np6 port-list
diagnose npu np6 port-list
```

Use either of the following command to list the NP6lite processors in your FortiGate unit:

```
get hardware npu np6lite port-list
diagnose npu np6lite port-list
```

To list other network processors on your FortiGate unit, use the following CLI command.

```
get hardware npu <model> list
<model> can be legacy, np1, np2 or np4.
```

The output lists the interfaces that have the specified processor. For example, for a FortiGate-5001B:

```
get hardware npu np4 list
ID      Model      Slot      Interface
0       On-board          port1 port2 port3 port4
        fabric1 base1 npu0-vlink0 npu0-vlink1
1       On-board          port5 port6 port7 port8
        fabric2 base2 npu1-vlink0 npu1-vlink1
```

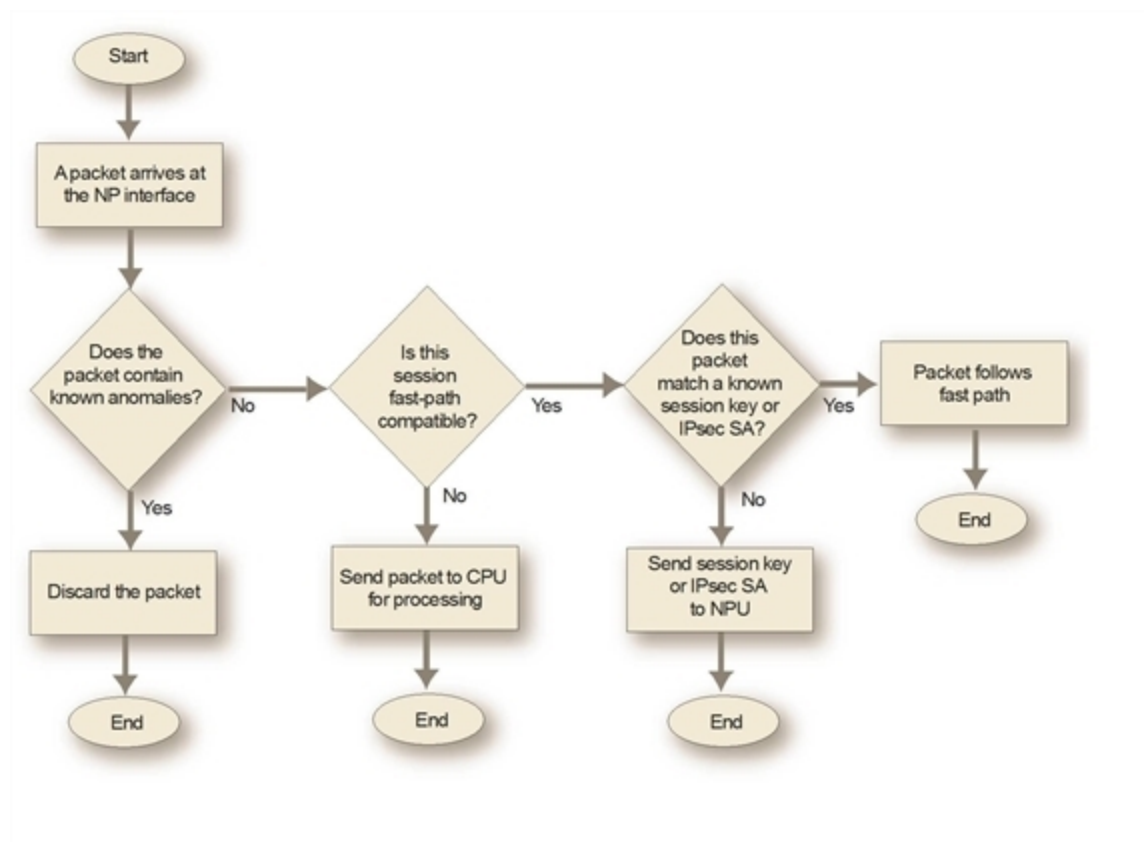
The npu0-vlink0, npu1-vlink1 etc interfaces are used for accelerating inter-VDOM links.

How NP hardware acceleration alters packet flow

NP hardware acceleration generally alters packet flow as follows:

1. Packets initiating a session pass to the FortiGate unit's main processing resources (CPU).
2. The FortiGate unit assesses whether the session matches fast path (offload) requirements. To be suitable for offloading, traffic must possess only characteristics that can be processed by the fast path. The list of requirements depends on the processor, see [NP6 session fast path requirements on page 1256](#) or [NP4 session fast path requirements on page 1329](#). If the session can be fast pathed, the FortiGate unit sends the session key or IPsec security association (SA) and configured firewall processing action to the appropriate network processor.
3. Network processors continuously match packets arriving on their attached ports against the session keys and SAs they have received.
 - If a network processor's network interface is configured to perform hardware accelerated anomaly checks, the network processor drops or accepts packets that match the configured anomaly patterns. These checks are separate from and in advance of anomaly checks performed by IPS, which is not compatible with network processor offloading. See [Offloading NP4 anomaly detection on page 1336](#).
 - The network processor next checks for a matching session key or SA. If a matching session key or SA is found, and if the packet meets packet requirements, the network processor processes the packet according to the configured action and then sends the resulting packet. This is the actual offloading step. Performing this processing on the NP processor improves overall performance because the NP processor is optimized for this task. As well, overall FortiGate performance is improved because the CPU has fewer sessions to process.

NP network processor packet flow



- If a matching session key or SA is not found, or if the packet does not meet packet requirements, the packet cannot be offloaded. The network processor sends the data to the FortiGate unit's CPU, which processes the packet.

Encryption and decryption of IPsec traffic originating from the FortiGate can utilize network processor encryption capabilities.

Packet forwarding rates vary by the percentage of offloadable processing and the type of network processing required by your configuration, but are independent of frame size. For optimal traffic types, network throughput can equal wire speed.

NP processors and traffic logging and monitoring

Except for the NP6, network processors do not count offloaded packets, and offloaded packets are not logged by traffic logging and are not included in traffic statistics and traffic log reports.

NP6 processors support per-session traffic and byte counters, Ethernet MIB matching, and reporting through messages resulting in traffic statistics and traffic log reporting.

Accelerated sessions on FortiView All Sessions page

When viewing sessions in the FortiView All Sessions console, NP4/ NP6 accelerated sessions are highlighted with an NP4 or NP6 icon. The tooltip for the icon includes the NP processor type and the total number of accelerated sessions.

You can also configure filtering to display FortiASIC sessions.

NP session offloading in HA active-active configuration

Network processors can improve network performance in active-active (load balancing) high availability (HA) configurations, even though traffic deviates from general offloading patterns, involving more than one network processor, each in a separate FortiGate unit. No additional offloading requirements apply.

Once the primary FortiGate unit's main processing resources send a session key to its network processor(s), network processor(s) on the primary unit can redirect any subsequent session traffic to other cluster members, reducing traffic redirection load on the primary unit's main processing resources.

As subordinate units receive redirected traffic, each network processor in the cluster assesses and processes session offloading independently from the primary unit. Session key states of each network processor are not part of synchronization traffic between HA members.

Configuring NP HMAC check offloading

Hash-based Message Authentication Code (HMAC) checks offloaded to network processors by default. You can enter the following command to disable this feature:

```
configure system global
    set ipsec-hmac-offload disable
end
```

Software switch interfaces and NP processors

FortiOS supports creating a software switch by grouping two or more FortiGate physical interfaces into a single virtual or software switch interface. All of the interfaces in this virtual switch act like interfaces in a hardware switch in that they all have the same IP address and can be connected to the same network. You create a software switch interface from the CLI using the command `config system switch-interface`.

The software switch is a bridge group of several interfaces, and the FortiGate CPU maintains the mac-port table for this bridge. As a result of this CPU involvement, traffic processed by a software switch interface is not offloaded to network processors.

Disabling NP acceleration for individual IPsec VPN phase 1s

Use the following command to disable NP offloading for an interface-based IPsec VPN phase 1:

```
config vpn ipsec phase1-interface
    edit phase-1-name
        set npu-offload disable
    end
```

Use the following command to disable NP offloading for a policy-based IPsec VPN phase 1:

```
config vpn ipsec phase1
    edit phase-1-name
        set npu-offload disable
    end
```

The `npu-offload` option is enabled by default.

Disabling NP offloading for unsupported IPsec encryption or authentication algorithms

In general, more recent IPsec VPN encryption and authentication algorithms may not be supported by older NP processors. For example, NP4 network processors do not support SHA-256, SHA-384, and SHA-512. IPsec traffic with unsupported algorithms is not offloaded and instead is processed by the FortiGate CPU. In addition, this

configuration may cause packet loss and other performance issues. If you experience packet loss or performance problems you should set the `npu-offload` option to `disable`. Future FortiOS versions should prevent selecting algorithms not supported by the hardware.

Disabling NP offloading for firewall policies

Use the following options to disable NP offloading for specific security policies:

For IPv4 security policies.

```
config firewall policy
  edit 1
    set auto-asic-offload disable
  end
```

For IPv6 security policies.

```
config firewall policy6
  edit 1
    set auto-asic-offload disable
  end
```

For multicast security policies.

```
config firewall multicast-policy
  edit 1
    set auto-asic-offload disable
  end
```

Enabling strict protocol header checking disables all hardware acceleration

You can use the following command to cause the FortiGate to apply strict header checking to verify that a packet is part of a session that should be processed. Strict header checking includes verifying the layer-4 protocol header length, the IP header length, the IP version, the IP checksum, IP options, and verifying that ESP packets have the correct sequence number, SPI, and data length. If the packet fails header checking it is dropped by the FortiGate unit.

```
config system global
  check-protocol-header strict
end
```

Enabling strict header checking disables all hardware acceleration. This includes NP, SP, and CP processing.

sFlow and NetFlow and hardware acceleration

NP6 offloading is supported when you configure NetFlow for interfaces connected to NP6 processors.

Configuring sFlow on any interface disables all NP4 and NP6 offloading for all traffic on that interface. As well, configuring NetFlow on any interface disables NP4 offloading for all traffic on that interface.

Checking that traffic is offloaded by NP processors

A number of diagnose commands can be used to verify that traffic is being offloaded.

Using the packet sniffer

Use the packet sniffer to verify that traffic is offloaded. Offloaded traffic is not picked up by the packet sniffer so if you are sending traffic through the FortiGate unit and it is not showing up on the packet sniffer you can conclude

that it is offloaded.

```
diag sniffer packet port1 <option>
```



If you want the packet sniffer to be able to see offloaded traffic you can temporarily disable offloading the traffic, run the packet sniffer to view it and then re-enable offloading. As an example, you may want to sniff the traffic that is accepted by a specific firewall policy. You can edit the policy and set the `auto-asic-offload` option to `disable` to disable offloading this traffic. You can also disable offloading for IPsec VPN traffic, see [Disabling NP acceleration for individual IPsec VPN phase 1s on page 1249](#).

Checking the firewall session offload tag

Use the `diagnose sys session list` command to display sessions. If the output for a session includes the `npu info` field you should see information about session being offloaded. If the output doesn't contain an `npu info` field then the session has not been offloaded.

```
diagnose sys session list
session info: proto=6 proto_state=01 duration=34 expire=3565 timeout=3600
              flags=00000000 sockflag=00000000 sockport=0 av_idx=0 use=3
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 policy_dir=0 tunnel=/
state=may_dirty npu
statistic(bytes/packets/allow_err): org=295/3/1 reply=60/1/1 tuples=2
orgin->sink: org pre->post, reply pre->post dev=48->6/6->48 gwy=10.1.100.11/11.11.11.1
hook=pre dir=org act=noop 172.16.200.55:56453->10.1.100.11:80(0.0.0.0:0)
hook=post dir=reply act=noop 10.1.100.11:80->172.16.200.55:56453(0.0.0.0:0)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=1 id_policy_id=0 auth_info=0 chk_client_info=0 vd=4
serial=0000091c tos=ff/ff ips_view=0 app_list=0 app=0
dd_type=0 dd_mode=0
per_ip_bandwidth meter: addr=172.16.200.55, bps=393
npu_state=00000000
npu info: flag=0x81/0x81, offload=4/4, ips_offload=0/0, epid=1/23, ipid=23/1,
vlan=32779/0
```

Verifying IPsec VPN traffic offloading

The following commands can be used to verify IPsec VPN traffic offloading to NP processors.

```
diagnose vpn ipsec status
NP1/NP2/NP4_0/sp_0_0:
  null: 0 0
  des: 0 0
    3des: 4075 4074
  aes: 0 0
  aria: 0 0
  seed: 0 0
  null: 0 0
    md5: 4075 4074
  sha1: 0 0
  sha256: 0 0
  sha384: 0 0
  sha512: 0 0
```

```

diagnose vpn tunnel list
list all ipsec tunnel in vd 3
-----
name=p1-vdom1 ver=1 serial=5 11.11.11.1:0->11.11.11.2:0 lgwy=static tun=tunnel
mode=auto bound_if=47
proxyid_num=1 child_num=0 refcnt=8 ilast=2 olast=2
stat: rxp=3076 txp=1667 rxb=4299623276 txb=66323
dpd: mode=active on=1 idle=5000ms retry=3 count=0 seqno=20
natt: mode=none draft=0 interval=0 remote_port=0
proxyid=p2-vdom1 proto=0 sa=1 ref=2 auto_negotiate=0 serial=1
src: 0:0.0.0.0/0.0.0.0:0
dst: 0:0.0.0.0/0.0.0.0:0
SA: ref=6 options=0000000e type=00 soft=0 mtu=1436 expire=1736 replaywin=2048 seqno=680
life: type=01 bytes=0/0 timeout=1748/1800
dec: spi=ae01010c esp=3des key=24 18e021bcace225347459189f292fbc2e4677563b07498a07
ah=md5 key=16 b4f44368741632b4e33e5f5b794253d3
enc: spi=ae01010d esp=3des key=24 42c94a8a2f72a44f9a3777f8e6aa3b24160b8af15f54a573
ah=md5 key=16 6214155f76b63a93345dcc9ec02d6415
dec:pkts/bytes=3073/4299621477, enc:pkts/bytes=1667/66375
  npu_flag=03 npu_rgwy=11.11.11.2 npu_lgwy=11.11.11.1 npu_selid=4

diagnose sys session list
session info: proto=6 proto_state=01 duration=34 expire=3565 timeout=3600
  flags=00000000 sockflag=00000000 sockport=0 av_idx=0 use=3
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 policy_dir=0 tunnel=/p1-vdom2
state=re may_dirty npu
statistic(bytes/packets/allow_err): org=112/2/1 reply=112/2/1 tuples=2
orgin->sink: org pre->post, reply pre->post dev=57->7/7->57 gwy=10.1.100.11/11.11.11.1
hook=pre dir=org act=noop 172.16.200.55:35254->10.1.100.11:80(0.0.0.0:0)
hook=post dir=reply act=noop 10.1.100.11:80->172.16.200.55:35254(0.0.0.0:0)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=1 id_policy_id=0 auth_info=0 chk_client_info=0 vd=4
serial=00002d29 tos=ff/ff ips_view=0 app_list=0 app=0
dd_type=0 dd_mode=0
per_ip_bandwidth meter: addr=172.16.200.55, bps=260
npu_state=00000000
npu info: flag=0x81/0x82, offload=7/7, ips_offload=0/0, epid=1/3, ipid=3/1,
vlan=32779/0

```

Dedicated management CPU

The web-based manager and CLI of FortiGate units with NP6 and NP4 processors may become unresponsive when the system is under heavy processing load because NP6 or NP4 interrupts overload the CPUs preventing CPU cycles from being used for management tasks. You can resolve this issue by using the following command to dedicate CPU core 0 to management tasks.

```

config system npu
  set dedicated-management-cpu {enable | disable}
end

```

All management tasks are then processed by CPU 0 and NP6 or NP4 interrupts are handled by the remaining CPU cores.

Offloading flow-based content inspection with NTurbo and IPSA

You can use the following command to configure NTurbo and IPA Advanced (IPSA) offloading and acceleration of firewall sessions that have flow-based security profiles. This includes firewall sessions with IPS, application control, CASI, flow-based antivirus and flow-based web filtering.

```
config ips global
  set np-accel-mode {none | basic}
  set cp-accel-mode {none | basic | advanced}
end
```

NTurbo offloads firewall sessions with flow-based security profiles to NPx processors

NTurbo offloads firewall sessions that include flow-based security profiles to NP4 or NP6 network processors. Without NTurbo, or with NTurbo disabled, all firewall sessions that include flow-based security profiles are processed by the FortiGate CPU.



NTurbo can only offload firewall sessions containing flow-based security profiles if the session could otherwise have been offloaded except for the presence of the flow-based security profiles. If something else prevents the session from being offloaded, NTurbo will not offload that session.



Firewall sessions that include proxy-based security profiles are never offloaded to network processors and are always processed by the FortiGate CPU.

NTurbo creates a special data path to redirect traffic from the ingress interface to IPS, and from IPS to the egress interface. NTurbo allows firewall operations to be offloaded along this path, and still allows IPS to behave as a stage in the processing pipeline, reducing the workload on the FortiGate CPU and improving overall throughput.



NTurbo sessions still offload pattern matching and other processes to CP processors, just like normal flow-based sessions.

If NTurbo is supported by your FortiGate unit, you can use the following command to configure it:

```
config ips global
  set np-accel-mode {basic | none}
end
```

`basic` enables NTurbo and is the default setting for FortiGate models that support NTurbo. `none` disables NTurbo. If the `np-accel-mode` option is not available, then your FortiGate does not support NTurbo.

There are some special cases where sessions may not be offloaded by NTurbo, even when NTurbo is explicitly enabled. In these cases the sessions are handled by the FortiGate CPU.

- NP acceleration is disabled. For example, `auto-asic-offload` is disabled in the firewall policy configuration.
- The firewall policy includes proxy-based security profiles.
- The sessions require FortiOS session-helpers. For example, FTP sessions can not be offloaded to NP processors because FTP sessions use the FTP session helper.

- Interface policies or DoS policies have been added to the ingress or egress interface.
- Tunneling is enabled. Any traffic to or from a tunneled interface (IPSec, IPinIP, SSL VPN, GRE, CAPWAP, etc.) cannot be offloaded by NTurbo.

IPSA offloads flow-based advanced pattern matching to CPx processors

IPSA offloads advanced or enhanced pattern matching operations required for flow-based content processing to CP8 and CP9 Content Processors. IPSA offloads enhanced pattern matching for NTurbo firewall sessions and firewall sessions that are not offloaded to NP processors. When IPSA is turned on, flow-based pattern databases are compiled and downloaded to the content processors from the IPS engine and IPS database. Flow-based pattern matching requests are redirected to the CP hardware reducing the load on the FortiGate CPU and accelerating pattern matching.

IF IPSA is supported on your FortiGate unit, you can use the following command to configure it:

```
config ips global
    set cp-accel-mode {advanced | basic | none}
end
```

`basic` offloads basic pattern matching. `advanced` offloads more types of pattern matching resulting in higher throughput than basic mode. `advanced` is only available on FortiGate models with two or more CP8s or one or more CP9s. If the `cp-accel-mode` option is not available, then your FortiGate does not support IPSA.

On FortiGates with one CP8, the default `cp-accel-mode` is `basic`. Setting the mode to `advanced` does not change the types of pattern matching that are offloaded.

On FortiGates with two or more CP8s or one or more CP9s the default `cp-accel-mode` is `advanced`. You can set the mode to `basic` to offload fewer types of pattern matching.

Preventing packet ordering problems with NP4, NP6 and NP6lite FortiGates under heavy load

In some cases when FortiGate units with NP4, NP6, or NP6lite processors are under heavy load the packets used in the TCP 3-way handshake of some sessions may be transmitted by the FortiGate in the wrong order resulting in the TCP sessions failing.

If you notice TCP sessions failing when a FortiGate with NP4, NP6, or NP6lite processors is very busy you can enable `delay-tcp-npu-session` in the firewall policy receiving the traffic. This option resolves the problem by delaying the session to make sure that there is time for all of the handshake packets to reach the destination before the session begins transmitting data.

```
config firewall policy
    set delay-tcp-npu-session enable
end
```

NP6 and NP6lite acceleration

NP6 and NP6lite network processors provide fastpath acceleration by offloading communication sessions from the FortiGate CPU. When the first packet of a new session is received by an interface connected to an NP6 processor, just like any session connecting with any FortiGate interface, the session is forwarded to the FortiGate CPU where it is matched with a security policy. If the session is accepted by a security policy and if the session can be offloaded its session key is copied to the NP6 processor that received the packet. All of the rest of the packets in the session are intercepted by the NP6 processor and fast-pathed out of the FortiGate unit to their destination without ever passing through the FortiGate CPU. The result is enhanced network performance provided by the NP6 processor plus the network processing load is removed from the CPU. In addition the NP6 processor can handle some CPU intensive tasks, like IPsec VPN encryption/decryption.



NP6lite processors have the same architecture and function in the same way as NP6 processors. All of the descriptions of NP6 processors in this document can be applied to NP6lite processors except where noted.

Session keys (and IPsec SA keys) are stored in the memory of the NP6 processor that is connected to the interface that received the packet that started the session. All sessions are fast-pathed and accelerated, even if they exit the FortiGate unit through an interface connected to another NP6. There is no dependence on getting the right pair of interfaces since the offloading is done by the receiving NP6.

The key to making this possible is an Integrated Switch Fabric (ISF) that connects the NP6s and the FortiGate unit interfaces together. Many FortiGate units with NP6 processors also have an ISF. The ISF allows any port connectivity. All ports and NP6s can communicate with each other over the ISF. There are no special ingress and egress fast path requirements as long as traffic enters and exits on interfaces connected to the same ISF.

Some FortiGate units, such as the FortiGate-1000D include multiple NP6 processors that are not connected by an ISF. Because the ISF is not present fast path acceleration is supported only between interfaces connected to the same NP6 processor. Since the ISF introduces some latency, models with no ISF provide low-latency network acceleration between network interfaces connected to the same NP6 processor.

Each NP6 has a maximum throughput of 40 Gbps using 4 x 10 Gbps XAUI or Quad Serial Gigabit Media Independent Interface (QSGMII) interfaces or 3 x 10 Gbps and 16 x 1 Gbps XAUI or QSGMII interfaces.

There are at least two limitations to keep in mind:

- The capacity of each NP6 processor. An individual NP6 processor can support between 10 and 16 million sessions. This number is limited by the amount of memory the processor has. Once an NP6 processor hits its session limit, sessions that are over the limit are sent to the CPU. You can avoid this problem by as much as possible distributing incoming sessions evenly among the NP6 processors. To be able to do this you need to be aware of which interfaces connect to which NP6 processors and distribute incoming traffic accordingly.

- The NP6 processors in some FortiGate units employ NP direct technology that removes the ISF. The result is very low latency but no inter-processor connectivity requiring you to make sure that traffic to be offloaded enters and exits the FortiGate through interfaces connected to the same NP processor.

NP6 session fast path requirements

NP6 processors can offload the following traffic and services:

- IPv4 and IPv6 traffic and NAT64 and NAT46 traffic (as well as IPv4 and IPv6 versions of the following traffic types where appropriate).
- Link aggregation (LAG) (IEEE 802.3ad) traffic (see [Increasing NP6 offloading capacity using link aggregation groups \(LAGs\) on page 1261](#)).
- TCP, UDP, ICMP, SCTP, and RDP traffic.
- IPsec VPN traffic, and offloading of IPsec encryption/decryption (including SHA2-256 and SHA2-512)
- IPsec traffic that passes through a FortiGate without being unencrypted.
- Anomaly-based intrusion prevention, checksum offload and packet defragmentation.
- IPIP tunneling (also called IP in IP tunneling), SIT tunneling, and IPv6 tunneling sessions.
- Multicast traffic (including Multicast over IPsec).
- CAPWAP and wireless bridge traffic tunnel encapsulation to enable line rate wireless forwarding from FortiAP devices (not supported by the NP6lite).
- Traffic shaping and priority queuing for both shared and per IP traffic shaping.
- Syn proxying (not supported by the NP6lite).
- DNS session helper (not supported by the NP6lite)/
- Inter-VDOM link traffic.

Sessions that are offloaded must be fast path ready. For a session to be fast path ready it must meet the following criteria:

- Layer 2 type/length must be 0x0800 for IPv4 or 0x86dd for IPv6 (IEEE 802.1q VLAN specification is supported).
- Layer 3 protocol can be IPv4 or IPv6.
- Layer 4 protocol can be UDP, TCP, ICMP, or SCTP.
- In most cases, Layer 3 / Layer 4 header or content modification sessions that require a session helper can be offloaded.
- Local host traffic (originated by the FortiGate unit) can be offloaded.
- If the FortiGate supports, NTurbo sessions can be offloaded if they are accepted by firewall policies that include IPS, Application Control, CASI, flow-based antivirus, or flow-based web filtering.

Offloading Application layer content modification is not supported. This means that sessions are not offloaded if they are accepted by firewall policies that include proxy-based virus scanning, proxy-based web filtering, DNS filtering, DLP, Anti-Spam, VoIP, ICAP, Web Application Firewall, or Proxy options.



If you disable anomaly checks by Intrusion Prevention (IPS), you can still enable hardware accelerated anomaly checks using the `fp-anomaly` field of the `config system interface` CLI command. See [Configuring individual NP6 processors on page 1265](#).

If a session is not fast path ready, the FortiGate unit will not send the session key or IPsec SA key to the NP6 processor. Without the session key, all session key lookup by a network processor for incoming packets of that

session fails, causing all session packets to be sent to the FortiGate unit's main processing resources, and processed at normal speeds.

If a session is fast path ready, the FortiGate unit will send the session key or IPsec SA key to the network processor. Session key or IPsec SA key lookups then succeed for subsequent packets from the known session or IPsec SA.

Packet fast path requirements

Packets within the session must then also meet packet requirements.

- Incoming packets must not be fragmented.
- Outgoing packets must not require fragmentation to a size less than 385 bytes. Because of this requirement, the configured MTU (Maximum Transmission Unit) for a network processor's network interfaces must also meet or exceed the network processors' supported minimum MTU of 385 bytes.

Mixing fast path and non-fast path traffic

If packet requirements are not met, an individual packet will be processed by the FortiGate CPU regardless of whether other packets in the session are offloaded to the NP6.

Also, in some cases, a protocol's session(s) may receive a mixture of offloaded and non-offloaded processing. For example, VoIP control packets may not be offloaded but VoIP data packets (voice packets) may be offloaded.

NP6Lite processors

The NP6Lite works the same way as the NP6. Being a lighter version, the NP6Lite has a lower capacity than the NP6. The NP6lite max throughput is 10 Gbps using 2x QSGMII and 2x Reduced gigabit media-independent interface (RGMII) interfaces.

Also, the NP6lite does not offload the following types of sessions:

- CAPWAP
- Syn proxy
- DNS session helper

NP6 and NP6Lite processors and sFlow and NetFlow

NP6 and NP6Lite offloading is supported when you configure NetFlow for interfaces connected to NP6 or NP6Lite processors. Offloading of other sessions is not affected by configuring NetFlow.

Configuring sFlow on any interface disables all NP6 and NP6Lite offloading for all traffic on that interface.

NP6 processors and traffic shaping

NP6-offloaded sessions support most types of traffic shaping. However, in bandwidth and out bandwidth traffic shaping, set using the following command, is not supported:

```
config system interface
  edit port1
    set outbandwidth <value>
    set inbandwidth <value>
  end
```


Configuring in bandwidth traffic shaping has no effect. Configuring out bandwidth traffic shaping imposes more limiting than configured, potentially reducing throughput more than expected.

NP Direct

On FortiGates with more than one NP6 processor, removing the Internal Switch Fabric (ISF) for NP Direct architecture provides direct access to the NP6 processors for the lowest latency forwarding. Because the NP6 processors are not connected, care must be taken with network design to make sure that all traffic to be offloaded enters and exits the FortiGate through interfaces connected to the same NP6 processor. As well Link Aggregation (LAG) interfaces should only include interfaces all connected to the same NP6 processor.

Example NP direct hardware with more than one NP6 processor includes:

- Ports 25 to 32 of the FortiGate-3700D in low latency mode.
- FortiGate-2000E
- FortiGate-2500E

Viewing your FortiGate NP6 processor configuration

Use either of the following commands to view the NP6 processor hardware configuration of your FortiGate unit:

```
get hardware npu np6 port-list
diagnose npu np6 port-list
```

If your FortiGate has NP6lite processors, you can use either of the following commands:

```
get hardware npu np6lite port-list
diagnose npu np6lite port-list
```

For example, for the FortiGate-5001D the output would be:

```
get hardware npu np6 port-list
```

Chip	XAUI	Ports	Max Speed	Cross-chip offloading
np6_0	0	port3	10G	Yes
	1			
	2	base1	1G	Yes
	3			
	0-3	port1	40G	Yes
	0-3	fabric1	40G	Yes
	0-3	fabric3	40G	Yes
	0-3	fabric5	40G	Yes
np6_1	0			
	1	port4	10G	Yes
	2			
	3	base2	1G	Yes
	0-3	port2	40G	Yes
	0-3	fabric2	40G	Yes
	0-3	fabric4	40G	Yes

For more example output for different FortiGate models, see [FortiGate NP6 architectures on page 1276](#) and [FortiGate NP6lite architectures on page 1326](#).

You can also use the following command to view the offloading features enabled or disabled on each of the NP6 processors in your FortiGate unit:

```
diagnose npu npu-feature
```

	np_0	np_1
Fastpath	Enabled	Enabled
Low-latency-mode	Disabled	Disabled
Low-latency-cap	No	No
IPv4 firewall	Yes	Yes
IPv6 firewall	Yes	Yes
IPv4 IPSec	Yes	Yes
IPv6 IPSec	Yes	Yes
IPv4 tunnel	Yes	Yes
IPv6 tunnel	Yes	Yes
GRE tunnel	No	No
IPv4 Multicast	Yes	Yes
IPv6 Multicast	Yes	Yes
CAPWAP	Yes	Yes

Disabling NP6 and NP6lite hardware acceleration (fastpath)

You can use the following command to disable NP6 offloading for all traffic. This option disables NP6 offloading for all traffic for all NP6 and NP6lite processors.

```
config system npu
    set fastpath disable
end
```

Optimizing NP6 performance by distributing traffic to XAUI links

On most FortiGate units with NP6 processors, the FortiGate interfaces are switch ports that connect to the NP6 processors with XAUI links. Packets pass from the interfaces to the NP6 processor over the XAUI links. Each NP6 processor has a 40 Gigabit bandwidth capacity. The four XAUI links each have a 10 Gigabit capacity for a total of 40 Gigabits.

On many FortiGate units with NP6 processors, the NP6 processors and the XAUI links are over-subscribed. Since the NP6 processors are connected by an Integrated Switch Fabric, you do not have control over how traffic is distributed to them. In fact traffic is distributed evenly by the ISF.

However, you can control how traffic is distributed to the XAUI links and you can optimize performance by distributing traffic evenly among the XAUI links. For example, if you have a very high amount of traffic passing between two networks, you can connect each network to interfaces connected to different XAUI links to distribute the traffic for each network to a different XAUI link.

For example, on a FortiGate-3200D (See [FortiGate-3200D fast path architecture on page 1300](#)), there are 48 10-Gigabit interfaces that send and receive traffic for two NP6 processors over a total of eight 10-Gigabit XAUI links. Each XAUI link gets traffic from six 10-Gigabit FortiGate interfaces. The amount of traffic that the FortiGate-3200D can offload is limited by the number of NP6 processors and the number of XAUI links. You can optimize the amount of traffic that the FortiGate-3200D can process by distributing it evenly among the XAUI links and the NP6 processors.

You can see the Ethernet interface, XAUI link, and NP6 configuration by entering the `get hardware npu np6 port-list` command. For the FortiGate-3200D the output is:

```

get hardware npu np6 port-list
Chip    XAUI Ports    Max    Cross-chip
        ----- Speed offloading
np6_0   0    port1    10G    Yes
        0    port5    10G    Yes
        0    port10   10G    Yes
        0    port13   10G    Yes
        0    port17   10G    Yes
        0    port22   10G    Yes
        1    port2    10G    Yes
        1    port6    10G    Yes
        1    port9    10G    Yes
        1    port14   10G    Yes
        1    port18   10G    Yes
        1    port21   10G    Yes
        2    port3    10G    Yes
        2    port7    10G    Yes
        2    port12   10G    Yes
        2    port15   10G    Yes
        2    port19   10G    Yes
        2    port24   10G    Yes
        3    port4    10G    Yes
        3    port8    10G    Yes
        3    port11   10G    Yes
        3    port16   10G    Yes
        3    port20   10G    Yes
        3    port23   10G    Yes
np6_1   0    port26   10G    Yes
        0    port29   10G    Yes
        0    port33   10G    Yes
        0    port37   10G    Yes
        0    port41   10G    Yes
        0    port45   10G    Yes
        1    port25   10G    Yes
        1    port30   10G    Yes
        1    port34   10G    Yes
        1    port38   10G    Yes
        1    port42   10G    Yes
        1    port46   10G    Yes
        2    port28   10G    Yes
        2    port31   10G    Yes
        2    port35   10G    Yes
        2    port39   10G    Yes
        2    port43   10G    Yes
        2    port47   10G    Yes
        3    port27   10G    Yes
        3    port32   10G    Yes
        3    port36   10G    Yes
        3    port40   10G    Yes
        3    port44   10G    Yes

```

```

      3      port48  10G      Yes
-----

```

In this command output you can see that each NP6 has for four XAUI links (0 to 3) and that each XAUI link is connected to six 10-gigabit Ethernet interfaces. To optimize throughput you should keep the amount of traffic being processed by each XAUI port to under 10 Gbps. So for example, if you want to offload traffic from four 10-gigabit networks you can connect these networks to Ethernet interfaces 1, 2, 3 and 4. This distributes the traffic from each 10-Gigabit network to a different XAUI link. Also, if you wanted to offload traffic from four more 10-Gigabit networks you could connect them to Ethernet ports 26, 25, 28, and 27. As a result each 10-Gigabit network would be connected to a different XAUI link.

Enabling bandwidth control between the ISF and NP6 XAUI ports

In some cases, the Internal Switch Fabric (ISF) buffer size may be larger than the buffer size of an NP6 XAUI port that receives traffic from the ISF. If this happens, burst traffic from the ISF may exceed the capacity of an XAUI port and sessions may be dropped.

You can use the following command to configure bandwidth control between the ISF and XAUI ports. Enabling bandwidth control can smooth burst traffic and keep the XAUI ports from getting overwhelmed and dropping sessions.

Use the following command to enable bandwidth control:

```

config system npu
    set sw-np-bandwidth {0G | 2G | 4G | 5G | 6G}
end

```

The default setting is 0G which means no bandwidth control. The other options limit the bandwidth to 2Gbps, 4Gbps and so on.

Increasing NP6 offloading capacity using link aggregation groups (LAGs)

NP6 processors can offload sessions received by interfaces in link aggregation groups (LAGs) (IEEE 802.3ad). A 802.3ad Link Aggregation and its management protocol, Link Aggregation Control Protocol (LACP) LAG combines more than one physical interface into a group that functions like a single interface with a higher capacity than a single physical interface. For example, you could use a LAG if you want to offload sessions on a 30 Gbps link by adding three 10-Gbps interfaces to the same LAG.

All offloaded traffic types are supported by LAGs, including IPsec VPN traffic. Just like with normal interfaces, traffic accepted by a LAG is offloaded by the NP6 processor connected to the interfaces in the LAG that receive the traffic to be offloaded. If all interfaces in a LAG are connected to the same NP6 processor, traffic received by that LAG is offloaded by that NP6 processor. The amount of traffic that can be offloaded is limited by the capacity of the NP6 processor.

If a FortiGate has two or more NP6 processors connected by an integrated switch fabric (ISF), you can use LAGs to increase offloading by sharing the traffic load across multiple NP6 processors. You do this by adding physical interfaces connected to different NP6 processors to the same LAG.

Adding a second NP6 processor to a LAG effectively doubles the offloading capacity of the LAG. Adding a third further increases offloading. The actual increase in offloading capacity may not actually be doubled by adding a second NP6 or tripled by adding a third. Traffic and load conditions and other factors may limit the actual offloading result.

The increase in offloading capacity offered by LAGs and multiple NP6s is supported by the integrated switch fabric (ISF) that allows multiple NP6 processors to share session information. Most FortiGate units with multiple NP6 processors also have an ISF. However, the FortiGate-1000D does not have an ISF. On this model and others that have more than one NP6 and no ISF, if you attempt to add interfaces connected to different NP6 processors to a LAG the system displays an error message.

There are also a few limitations to LAG NP6 offloading support for IPsec VPN:

- IPsec VPN anti-replay protection cannot be used if IPsec is configured on a LAG that has interfaces connected to multiple NP6 processors.
- Because the encrypted traffic for one IPsec VPN tunnel has the same 5-tuple, the traffic from one tunnel can only be balanced to one interface in a LAG. This limits the maximum throughput for one IPsec VPN tunnel in an NP6 LAG group to 10Gbps.

Configuring inter-VDOM link acceleration with NP6 processors

FortiGate units with NP6 processors include inter-VDOM links that can be used to accelerate inter-VDOM link traffic.

- For a FortiGate unit with two NP6 processors there are two accelerated inter-VDOM links, each with two interfaces:
 - **npu0_vlink:**
npu0_vlink0
npu0_vlink1
 - **npu1_vlink:**
npu1_vlink0
npu1_vlink1

These interfaces are visible from the GUI and CLI. For a FortiGate unit with NP6 interfaces, enter the following CLI command to display the NP6-accelerated inter-VDOM links:

```
get system interface
...
== [ npu0_vlink0 ]
name: npu0_vlink0 mode: static ip: 0.0.0.0 0.0.0.0 status: down netbios-forward: disable
type: physical sflow-sampler: disable explicit-web-proxy: disable explicit-ftp-proxy:
disable mtu-override: disable wccp: disable drop-overlapped-fragment: disable drop-
fragment: disable

== [ npu0_vlink1 ]
name: npu0_vlink1 mode: static ip: 0.0.0.0 0.0.0.0 status: down netbios-forward: disable
type: physical sflow-sampler: disable explicit-web-proxy: disable explicit-ftp-proxy:
disable mtu-override: disable wccp: disable drop-overlapped-fragment: disable drop-
fragment: disable

== [ npu1_vlink0 ]
name: npu1_vlink0 mode: static ip: 0.0.0.0 0.0.0.0 status: down netbios-forward: disable
type: physical sflow-sampler: disable explicit-web-proxy: disable explicit-ftp-proxy:
disable mtu-override: disable wccp: disable drop-overlapped-fragment: disable drop-
fragment: disable

== [ npu1_vlink1 ]
name: npu1_vlink1 mode: static ip: 0.0.0.0 0.0.0.0 status: down netbios-forward: disable
type: physical sflow-sampler: disable explicit-web-proxy: disable explicit-ftp-proxy:
disable mtu-override: disable wccp: disable drop-overlapped-fragment: disable drop-
fragment: disable
...
```

By default the interfaces in each inter-VDOM link are assigned to the root VDOM. To use these interfaces to accelerate inter-VDOM link traffic, assign each interface in the pair to the VDOMs that you want to offload traffic between. For example, if you have added a VDOM named New-VDOM to a FortiGate unit with NP4 processors, you can go to **System > Network > Interfaces** and edit the **npu0-vlink1** interface and set the **Virtual Domain** to **New-VDOM**. This results in an accelerated inter-VDOM link between root and New-VDOM. You can also do this from the CLI:

```
config system interface
  edit npu0-vlink1
    set vdom New-VDOM
  end
```

Using VLANs to add more accelerated inter-VDOM links

You can add VLAN interfaces to the accelerated inter-VDOM links to create inter-VDOM links between more VDOMs. For the links to work, the VLAN interfaces must be added to the same inter-VDOM link, must be on the same subnet, and must have the same VLAN ID.

For example, to accelerate inter-VDOM link traffic between VDOMs named Marketing and Engineering using VLANs with VLAN ID 100 go to **System > Network > Interfaces** and select **Create New** to create the VLAN interface associated with the Marketing VDOM:

Name	Marketing-link
Type	VLAN
Interface	npu0_vlink0
VLAN ID	100
Virtual Domain	Marketing
IP/Network Mask	172.20.120.12/24

Create the inter-VDOM link associated with Engineering VDOM:

Name	Engineering-link
Type	VLAN
Interface	npu0_vlink1
VLAN ID	100
Virtual Domain	Engineering
IP/Network Mask	172.20.120.22/24

Or do the same from the CLI:

```
config system interface
  edit Marketing-link
    set vdom Marketing
```

```

set ip 172.20.120.12/24
set interface npu0_vlink0
set vlanid 100
next
edit Engineering-link
set vdom Engineering
set ip 172.20.120.22/24
set interface npu0_vlink1
set vlanid 100

```

Confirm that the traffic is accelerated

Use the following CLI commands to obtain the interface index and then correlate them with the session entries. In the following example traffic was flowing between new accelerated inter-VDOM links and physical ports port1 and port 2 also attached to the NP6 processor.

diagnose ip address list

```

IP=172.31.17.76->172.31.17.76/255.255.252.0 index=5 devname=port1
IP=10.74.1.76->10.74.1.76/255.255.252.0 index=6 devname=port2
IP=172.20.120.12->172.20.120.12/255.255.255.0 index=55 devname=IVL-VLAN1_ROOT
IP=172.20.120.22->172.20.120.22/255.255.255.0 index=56 devname=IVL-VLAN1_VDOM1

```

diagnose sys session list

```

session info: proto=1 proto_state=00 duration=282 expire=24 timeout=0 session info:
    proto=1 proto_state=00 duration=124 expire=59 timeout=0 flags=00000000
    sockflag=00000000 sockport=0 av_idx=0 use=3
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 policy_dir=0 tunnel=/
state=may_dirty npu
statistic(bytes/packets/allow_err): org=180/3/1 reply=120/2/1 tuples=2
origin->sink: org pre->post, reply pre->post dev=55->5/5->55
    gwy=172.31.19.254/172.20.120.22
hook=post dir=org act=snat 10.74.2.87:768->10.2.2.2:8(172.31.17.76:62464)
hook=pre dir=reply act=dnat 10.2.2.2:62464->172.31.17.76:0(10.74.2.87:768)
misc=0 policy_id=4 id_policy_id=0 auth_info=0 chk_client_info=0 vd=0
serial=0000004e tos=ff/ff ips_view=0 app_list=0 app=0
dd_type=0 dd_mode=0
per_ip_bandwidth meter: addr=10.74.2.87, bps=880
npu_state=00000000
npu info: flag=0x81/0x81, offload=8/8, ips_offload=0/0, epid=160/218, ipid=218/160,
vlan=32769/0

```

```

session info: proto=1 proto_state=00 duration=124 expire=20 timeout=0 flags=00000000
    sockflag=00000000 sockport=0 av_idx=0 use=3
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 policy_dir=0 tunnel=/
state=may_dirty npu
statistic(bytes/packets/allow_err): org=180/3/1 reply=120/2/1 tuples=2
origin->sink: org pre->post, reply pre->post dev=6->56/56->6
    gwy=172.20.120.12/10.74.2.87
hook=pre dir=org act=noop 10.74.2.87:768->10.2.2.2:8(0.0.0.0:0)
hook=post dir=reply act=noop 10.2.2.2:768->10.74.2.87:0(0.0.0.0:0)
misc=0 policy_id=3 id_policy_id=0 auth_info=0 chk_client_info=0 vd=1

```

```

serial=0000004d tos=ff/ff ips_view=0 app_list=0 app=0
dd_type=0 dd_mode=0
per_ip_bandwidth meter: addr=10.74.2.87, bps=880
npu_state=00000000
npu info: flag=0x81/0x81, offload=8/8, ips_offload=0/0, epid=219/161, ipid=161/219,
          vlan=0/32769
total session 2

```

Disabling offloading IPsec Diffie-Hellman key exchange

You can use the following command to disable using ASIC offloading to accelerate IPsec Diffie-Hellman key exchange for IPsec ESP traffic. By default hardware offloading is used. For debugging purposes or other reasons you may want this function to be processed by software.

Use the following command to disable using ASIC offloading for IPsec Diffie-Hellman key exchange:

```

config system global
    set ipsec-asic-offload disable
end

```

Configuring individual NP6 processors

You can use the `config system np6` command to configure a wide range of settings for each of the NP6 processors in your FortiGate unit including enabling session accounting and adjusting session timeouts. As well you can set anomaly checking for IPv4 and IPv6 traffic.

You can also enable and adjust Host Protection Engine (HPE) to protect networks from DoS attacks by categorizing incoming packets based on packet rate and processing cost and applying packet shaping to packets that can cause DoS attacks.

The settings that you configure for an NP6 processor with the `config system np6` command apply to traffic processed by all interfaces connected to that NP6 processor. This includes the physical interfaces connected to the NP6 processor as well as all subinterfaces, VLAN interfaces, IPsec interfaces, LAGs and so on associated with the physical interfaces connected to the NP6 processor.



Some of the options for this command apply anomaly checking for NP6 sessions in the same way as the command described in [Offloading NP4 anomaly detection on page 1336](#) applies anomaly checking for NP4 sessions.

```

config system np6
    edit <np6-processor-name>
        set low-latency-mode {disable | enable}
        set per-session-accounting {all-enable | disable | enable-by-log}
        set session-timeout-random-range <range>
        set garbage-session-collector {disable | enable}
        set session-collector-interval <range>
        set session-timeout-interval <range>
        set session-timeout-random-range <range>
        set session-timeout-fixed {disable | enable}
        config hpe
            set tcpsyn-max <packets-per-second>
            set tcp-max <packets-per-second>
            set udp-max <packets-per-second>
            set icmp-max <packets-per-second>
            set sctp-max <packets-per-second>
        end
    end
end

```



```

set esp-max <packets-per-second>
set ip-frag-max <packets-per-second>
set ip-others-max <packets-per-second>
set arp-max <packets-per-second>
set l2-others-max <packets-per-second>
set enable-shaper {disable | enable}
config fp-anomaly-v4
set tcp-syn-fin {allow | drop | trap-to-host}
set tcp_fin_noack {allow | drop | trap-to-host}
set tcp_fin_only {allow | drop | trap-to-host}
set tcp_no_flag {allow | drop | trap-to-host}
set tcp_syn_data {allow | drop | trap-to-host}
set tcp-winnuke {allow | drop | trap-to-host}
set tcp-land {allow | drop | trap-to-host}
set udp-land {allow | drop | trap-to-host}
set icmp-land {allow | drop | trap-to-host}
set icmp-frag {allow | drop | trap-to-host}
set ipv4-land {allow | drop | trap-to-host}
set ipv4-proto-err {allow | drop | trap-to-host}
set ipv4-unknopt {allow | drop | trap-to-host}
set ipv4-optrr {allow | drop | trap-to-host}
set ipv4-optssrr {allow | drop | trap-to-host}
set ipv4-optlsrr {allow | drop | trap-to-host}
set ipv4-optstream {allow | drop | trap-to-host}
set ipv4-optsecurity {allow | drop | trap-to-host}
set ipv4-opttimestamp {allow | drop | trap-to-host}
set ipv4-csum-err {drop | trap-to-host}
set tcp-csum-err {drop | trap-to-host}
set udp-csum-err {drop | trap-to-host}
set icmp-csum-err {drop | trap-to-host}
set ipv6-land {allow | drop | trap-to-host}
set ipv6-proto-err {allow | drop | trap-to-host}
set ipv6-unknopt {allow | drop | trap-to-host}
set ipv6-saddr-err {allow | drop | trap-to-host}
set ipv6-daddr-err {allow | drop | trap-to-host}
set ipv6-optralert {allow | drop | trap-to-host}
set ipv6-optjumbo {allow | drop | trap-to-host}
set ipv6-opttunnel {allow | drop | trap-to-host}
set ipv6-opthomeaddr {allow | drop | trap-to-host}
set ipv6-optnsap {allow | drop | trap-to-host}
set ipv6-optendpid {allow | drop | trap-to-host}
set ipv6-optinvld {allow | drop | trap-to-host}
end

```

Command syntax

Command	Description	Default
low-latency-mode {disable enable}	Enable low-latency mode. In low latency mode the integrated switch fabric is bypassed. Low latency mode requires that packet enter and exit using the same NP6 processor. This option is only available for NP6 processors that can operate in low-latency mode, currently only np6_0 and np6_1 on the FortiGate-3700D and DX.	disable

Command	Description	Default
<code>per-session-accounting</code> {all-enable disable enable-by-log}	Disable NP6 per-session accounting or enable it and control how it works. If set to <code>enable-by-log</code> (the default) NP6 per-session accounting is only enabled if firewall policies accepting offloaded traffic have traffic logging enabled. If set to <code>all-enable</code> , NP6 per-session accounting is always enabled for all traffic offloaded by the NP6 processor. Enabling per-session accounting can affect performance.	enable-by-log
<code>garbage-session-collector</code> {disable enable}	Enable deleting expired or garbage sessions.	disable
<code>session-collector-interval</code> <range>	Set the expired or garbage session collector time interval in seconds. The range is 1 to 100 seconds.	64
<code>session-timeout-interval</code> <range>	Set the timeout for checking for and removing inactive NP6 sessions. The range is 0 to 1000 seconds.	40
<code>session-timeout-random-range</code> <range>	Set the random timeout for checking and removing inactive NP6 sessions. The range is 0 to 1000 seconds.	8
<code>session-timeout-fixed</code> {disable enable}	Enable to force checking for and removing inactive NP6 sessions at the <code>session-timeout-interval</code> time interval. Set to <code>disable</code> (the default) to check for and remove inactive NP6 sessions at random time intervals.	disable
config hpe		
<code>hpe</code>	Use the following options to use HPE to apply DDoS protection at the NP6 processor by limiting the number packets per second received for various packet types by each NP6 processor. This rate limiting is applied very efficiently because it is done in hardware by the NP6 processor.	
<code>enable-shaper</code> {disable enable}	Enable or disable HPE DDoS protection.	disable
<code>tcpsyn-max</code>	Limit the maximum number of TCP SYN packets received per second. The range is 10,000 to 4,000,000,000 pps. The default limits the number of packets per second to 5,000,000 pps.	5000000

Command	Description	Default
<code>tcp-max</code>	Limit the maximum number of TCP packets received per second. The range is 10,000 to 4,000,000,000 pps. The default limits the number of packets per second to 5,000,000 pps.	5000000
<code>udp-max</code>	Limit the maximum number of UDP packets received per second. The range is 10,000 to 4,000,000,000 pps. The default limits the number of packets per second to 5,000,000 pps.	5000000
<code>icmp-max</code>	Limit the maximum number of ICMP packets received. The range is 10,000 to 4,000,000,000 pps. The default is 100,000 pps.	100000
<code>sctp-max</code>	Limit the maximum number of SCTP packets received. The range is 10,000 to 4,000,000,000 pps. The default is 100,000 pps.	100000
<code>esp-max</code>	Limit the maximum number of ESP packets received. The range is 10,000 to 4,000,000,000 pps. The default is 100,000 pps.	100000
<code>ip-frag-max</code>	Limit the maximum number of fragmented IP packets received. The range is 10,000 to 4,000,000,000 pps. The default is 100,000 pps.	100000
<code>ip-others-max</code>	Limit the maximum number of other types of IP packets received. The range is 10,000 to 4,000,000,000 pps. The default is 100,000 pps.	100000
<code>arp-max</code>	Limit the maximum number of ARP packets received. The range is 10,000 to 4,000,000,000 pps. The default is 100,000 pps.	100000
<code>l2-others-max</code>	Limit the maximum number of other layer-2 packets received. The range is 10,000 to 4,000,000,000 pps. The default is 100,000 pps.	100000
config fp-anomaly		
<code>fp-anomaly-v4</code>	Configure how the NP6 processor does traffic anomaly protection. In most cases you can configure the NP6 processor to allow or drop the packets associated with an attack or forward the packets that are associated with the attack to FortiOS (called <code>trap-to-host</code>). Selecting <code>trap-to-host</code> turns off NP6 anomaly protection for that anomaly. If you require anomaly protection but don't want to use the NP6 processor, you can select <code>trap-to-host</code> and enable anomaly protection with a DoS policy.	


Command	Description	Default
<code>tcp-syn-fin {allow drop trap-to-host}</code>	Detects TCP SYN flood SYN/FIN flag set anomalies.	allow
<code>tcp_fin_noack {allow drop trap-to-host}</code>	Detects TCP SYN flood with FIN flag set without ACK setting anomalies.	trap-to-host
<code>tcp_fin_only {allow drop trap-to-host}</code>	Detects TCP SYN flood with only FIN flag set anomalies.	trap-to-host
<code>tcp_no_flag {allow drop trap-to-host}</code>	Detects TCP SYN flood with no flag set anomalies.	allow
<code>tcp_syn_data {allow drop trap-to-host}</code>	Detects TCP SYN flood packets with data anomalies.	allow
<code>tcp-winnuke {allow drop trap-to-host}</code>	Detects TCP WinNuke anomalies.	trap-to-host
<code>tcp-land {allow drop trap-to-host}</code>	Detects TCP land anomalies.	trap-to-host
<code>udp-land {allow drop trap-to-host}</code>	Detects UDP land anomalies.	trap-to-host
<code>icmp-land {allow drop trap-to-host}</code>	Detects ICMP land anomalies.	trap-to-host
<code>icmp-frag {allow drop trap-to-host}</code>	Detects Layer 3 fragmented packets that could be part of a layer 4 ICMP anomalies.	allow
<code>ipv4-land {allow drop trap-to-host}</code>	Detects IPv4 land anomalies.	trap-to-host
<code>ipv4-proto-err {allow drop trap-to-host}</code>	Detects invalid layer 4 protocol anomalies.	trap-to-host
<code>ipv4-unknopt {allow drop trap-to-host}</code>	Detects unknown option anomalies.	trap-to-host
<code>ipv4-optrr {allow drop trap-to-host}</code>	Detects IPv4 with record route option anomalies.	trap-to-host
<code>ipv4-optssrr {allow drop trap-to-host}</code>	Detects IPv4 with strict source record route option anomalies.	trap-to-host
<code>ipv4-optlsrr {allow drop trap-to-host}</code>	Detects IPv4 with loose source record route option anomalies.	trap-to-host

Command	Description	Default
ipv4-optstream {allow drop trap-to-host}	Detects stream option anomalies.	trap-to-host
ipv4-optsecurity {allow drop trap-to-host}	Detects security option anomalies.	trap-to-host
ipv4-opttimestamp {allow drop trap-to-host}	Detects timestamp option anomalies.	trap-to-host
ipv4-csum-err {drop trap-to-host}	Detects IPv4 checksum errors.	drop
tcp-csum-err {drop trap-to-host}	Detects TCP checksum errors.	drop
udp-csum-err {drop trap-to-host}	Detects UDP checksum errors.	drop
icmp-csum-err {drop trap-to-host}	Detects ICMP checksum errors.	drop
ipv6-land {allow drop trap-to-host}	Detects IPv6 land anomalies	trap-to-host
ipv6-unknopt {allow drop trap-to-host}	Detects unknown option anomalies.	trap-to-host
ipv6-saddr-err {allow drop trap-to-host}	Detects source address as multicast anomalies.	trap-to-host
ipv6-daddr_err {allow drop trap-to-host}	Detects destination address as unspecified or loopback address anomalies.	trap-to-host
ipv6-optralert {allow drop trap-to-host}	Detects router alert option anomalies.	trap-to-host
ipv6-optjumbo {allow drop trap-to-host}	Detects jumbo options anomalies.	trap-to-host
ipv6-opttunnel {allow drop trap-to-host}	Detects tunnel encapsulation limit option anomalies.	trap-to-host
ipv6-opthomeaddr {allow drop trap-to-host}	Detects home address option anomalies.	trap-to-host
ipv6-optnsap {allow drop trap-to-host}	Detects network service access point address option anomalies.	trap-to-host

Command	Description	Default
<code>ipv6-optendpid {allow drop trap-to-host}</code>	Detects end point identification anomalies.	trap-to-host
<code>ipv6-optinvld {allow drop trap-to-host}</code>	Detects invalid option anomalies.	trap-to-host

Enabling per-session accounting for offloaded NP6 and NP6lite sessions

Per-session accounting is a logging feature that allows the FortiGate to report the correct bytes/pkt numbers per session for sessions offloaded to an NP6 or NP6lite processor. This information appears in traffic log messages as well as in FortiView. The following example shows the Sessions dashboard widget tracking SPU and nTurbo sessions.



You can hover over the SPU icon to see some information about the offloaded sessions.

You configure per-session accounting for each NP6 processor. For example, use the following command to enable per-session accounting for NP6_0 and NP6_1:

```
config system np6
  edit np6_0
    set per-session-accounting enable-by-log
  next
  edit np6_1
    set per-session-accounting enable-by-log
  end
```

If your FortiGate has NP6lite processors, you can use the following command to enable per-session accounting for all of the NP6lite processors in the FortiGate unit:

```
config system npu
  set per-session-accounting enable-by-log
end
```

The option, `enable-by-log` enables per-session accounting for offloaded sessions with traffic logging enabled and `all-enable` enables per-session accounting for all offloaded sessions.

By default, `per-session-accounting` is set to `enable-by-log`, which results in per-session accounting being turned on when you enable traffic logging in a policy.

Per-session accounting can affect offloading performance. So you should only enable per-session accounting if you need the accounting information.

Enabling per-session accounting does not provide traffic flow data for sFlow or NetFlow.

Configuring NP6 session timeouts

For NP6 traffic, FortiOS refreshes an NP6 session's lifetime when it receives a session update message from the NP6 processor. To avoid session update message congestion, these NP6 session checks are performed all at once after a random time interval and all of the update messages are sent from the NP6 processor to FortiOS at once. This can result in fewer messages being sent because they are only sent at random time intervals instead of every time a session times out.

In fact, if your NP6 processor is processing a lot of short lived sessions, it is recommended that you use the default setting of random checking every 8 seconds to avoid very bursty session updates. If the time between session updates is very long and very many sessions have been expired between updates a large number of updates will need to be done all at once.

You can use the following command to set the random time range.

```
config system np6
  edit <np6-processor-name>
    set session-timeout-fixed disable
    set session-timeout-random-range 8
  end
```

This is the default configuration. The random timeout range is 1 to 1000 seconds and the default range is 8. So, by default, NP6 sessions are checked at random time intervals of between 1 and 8 seconds. So sessions can be inactive for up to 8 seconds before they are removed from the FortiOS session table.

If you want to reduce the amount of checking you can increase the `session-timeout-random-range`. This could result in inactive sessions being kept in the session table longer. But if most of your NP6 sessions are relatively long this shouldn't be a problem.

You can also change this session checking to a fixed time interval and set a fixed timeout:

```
config system np6
  edit <np6-processor-name>
    set session-timeout-fixed enable
    set session-timeout-fixed 40
  end
```

The fixed timeout default is every 40 seconds and the range is 1 to 1000 seconds. Using a fixed interval further reduces the amount of checking that occurs.

You can select random or fixed updates and adjust the time intervals to minimize the refreshing that occurs while still making sure inactive sessions are deleted regularly. For example, if an NP6 processor is processing sessions with long lifetimes you can reduce checking by setting a relatively long fixed timeout.

Configure the number of IPsec engines NP6 processors use

NP6 processors use multiple IPsec engines to accelerate IPsec encryption and decryption. In some cases out of order ESP packets can cause problems if multiple IPsec engines are running. To resolve this problem you can configure all of the NP6 processors to use fewer IPsec engines.

Use the following command to change the number of IPsec engines used for decryption (`ipsec-dec-subengine-mask`) and encryption (`ipsec-enc-subengine-mask`). These settings are applied to all of the NP6 processors in the FortiGate unit.

```
config system npu
    set ipsec-dec-subengine-mask <engine-mask>
    set ipsec-enc-subengine-mask <engine-mask>
end
```

<engine-mask> is a hexadecimal number in the range 0x01 to 0xff where each bit represents one IPsec engine. The default <engine-mask> for both options is 0xff which means all IPsec engines are used. Add a lower <engine-mask> to use fewer engines. You can configure different engine masks for encryption and decryption.

Stripping clear text padding and IPsec session ESP padding

In some situations, when clear text or ESP packets in IPsec sessions may have large amounts of layer 2 padding, the NP6 IPsec engine may not be able to process them and the session may be blocked.

If you notice dropped IPsec sessions, you could try using the following CLI options to cause the NP6 processor to strip clear text padding and ESP padding before send the packets to the IPsec engine. With padding stripped, the session can be processed normally by the IPsec engine.

Use the following command to strip ESP padding:

```
config system npu
    set strip-esp-padding enable
    set strip-clear-text-padding enable
end
```

Stripping clear text and ESP padding are both disabled by default.

Disable NP6 CAPWAP offloading

By default and where possible, managed FortiAP and FortiLink CAPWAP sessions are offloaded to NP6 processors. You can use the following command to disable CAPWAP session offloading:

```
config system npu
    set capwap-offload disable
end
```

Optionally disable NP6 offloading of traffic passing between 10Gbps and 1Gbps interfaces

Due to NP6 internal packet buffer limitations, some offloaded packets received at a 10Gbps interface and destined for a 1Gbps interface can be dropped, reducing performance for TCP and IP tunnel traffic. If you experience this performance reduction, you can use the following command to disable offloading sessions passing from 10Gbps interfaces to 1Gbps interfaces:

```
config system npu
    set host-shortcut-mode host-shortcut
end
```

Select `host-shortcut` to stop offloading TCP and IP tunnel packets passing from 10Gbps interfaces to 1Gbps interfaces. TCP and IP tunnel packets passing from 1Gbps interfaces to 10Gbps interfaces are still offloaded as normal.

If `host-shortcut` is set to the default `bi-directional` setting, packets in both directions are offloaded.

This option is only available if your FortiGate has 10G and 1G interfaces accelerated by NP6 processors.

Offloading RDP traffic

FortiOS supports NP6 offloading of Reliable Data Protocol (RDP) traffic. RDP is a network transport protocol that optimizes remote loading, debugging, and bulk transfer of images and data. RDP traffic uses Assigned Internet Protocol number 27 and is defined in [RFC 908](#) and updated in [RFC 1151](#). If your network is processing a lot of RDP traffic, offloading it can improve overall network performance.

You can use the following command to enable or disable NP6 RDP offloading. RDP offloading is enabled by default.

```
config system npu
    set rdp-offload {disable | enable}
end
```

NP6 session drift

In some cases, sessions processed by NP6 processors may fail to be deleted leading to a large number of idle sessions. This is called session drift. You can use SNMP to be alerted when the number of idle sessions becomes high. SNMP also allows you to see which NP6 processor has the abnormal number of idle sessions and you can use a diagnose command to delete them.

The following MIB fields allow you to use SNMP to monitor session table information for NP6 processors including drift for each NP6 processor:

```
FORTINET-FORTIGATE-MIB::fgNPUNumber.0 = INTEGER: 2
FORTINET-FORTIGATE-MIB::fgNPUName.0 = STRING: NP6
FORTINET-FORTIGATE-MIB::fgNPUDrvDriftSum.0 = INTEGER: 0
FORTINET-FORTIGATE-MIB::fgNPUIndex.0 = INTEGER: 0
FORTINET-FORTIGATE-MIB::fgNPUIndex.1 = INTEGER: 1
FORTINET-FORTIGATE-MIB::fgNPUSessionTblSize.0 = Gauge32: 33554432
FORTINET-FORTIGATE-MIB::fgNPUSessionTblSize.1 = Gauge32: 33554432
FORTINET-FORTIGATE-MIB::fgNPUSessionCount.0 = Gauge32: 0
FORTINET-FORTIGATE-MIB::fgNPUSessionCount.1 = Gauge32: 0
FORTINET-FORTIGATE-MIB::fgNPUDrvDrift.0 = INTEGER: 0
FORTINET-FORTIGATE-MIB::fgNPUDrvDrift.1 = INTEGER: 0
```

You can also use the following diagnose command to determine if drift is occurring:

```
diagnose npu np6 sse-drift-summary
NPU    drv-drift
-----
np6_0  0
np6_1  0
-----
Sum    0
-----
```

The command output shows a drift summary for all the NP6 processors in the system, and shows the total drift. Normally the sum is 0. The previous command output, from a FortiGate-1500D, shows that the 1500D's two NP6 processors are not experiencing any drift.

If the sum is not zero, then extra idle sessions may be accumulating. You can use the following command to delete those sessions:

```
diagnose npu np6 sse-purge-drift <np6_id> [<time>]
```

Where `<np6_id>` is the number (starting with NP6_0 with a `np6_id` of 0) of the NP6 processor for which to delete idle sessions in. `<time>` is the age in seconds of the idle sessions to be deleted. All idle sessions this age and older are deleted. The default time is 300 seconds.

The `diagnose npu np6 sse-stats <np6_id>` command output also includes a `drv-drift` field that shows the total drift for one NP6 processor.

Optimizing FortiGate-3960E and 3980E IPsec VPN performance

You can use the following command to configure outbound hashing to improve IPsec VPN performance for the FortiGate-3960E and 3980E. If you change these settings, to make sure they take affect, you should reboot your device.

```
config system np6
  edit np6_0
    set ipsec-outbound-hash {disable | enable}
    set ipsec-ob-hash-function {switch-group-hash | global- hash | global-hash-weighted
                              | round-robin-switch-group | round-robin-global}
  end
```

Where:

`ipsec-outbound-hash` is disabled by default. If you enable it you can set `ipsec-ob-hash-function` as follows:

`switch-group-hash` (the default) distribute outbound IPsec Security Association (SA) traffic to NP6 processors connected to the same switch as the interfaces that received the incoming traffic. This option, keeps all traffic on one switch and the NP6 processors connected to that switch, to improve performance.

`global-hash` distribute outbound IPsec SA traffic among all NP6 processors.

`global-hash-weighted` distribute outbound IPsec SA traffic from switch 1 among all NP6 processors with more sessions going to the NP6s connected to switch 0. This options is only recommended for the FortiGate-3980E because it is designed to weigh switch 0 hider to send more sessions to switch 0 which on the FortiGate-3980E has more NP6 processors connected to it. On the FortiGate-3960E both switches have the same number of NP6s so for best performance one switch shouldn't have a higher weight.

`round-robin-switch-group` round-robin distribution of outbound IPsec SA traffic among the NP6 processors connected to the same switch.

`round-robin-global` round-robin distribution of outbound IPsec SA traffic among all NP6 processors.

Recalculating packet checksums if the `iph.reserved` bit is set to 0

NP6 and the NP6lite processors clear the `iph.flags.reserved` bit. This results in the packet checksum becoming incorrect because by default the packet is changed but the checksum is not recalculated. Since the checksum is incorrect these packets may be dropped by the network stack. You can enable this option to cause the system to re-calculate the checksum. Enabling this option may cause a minor performance reduction. This option is disabled by default.

To enabled checksum recalculation for packets with the `iph.flags.reserved` header:

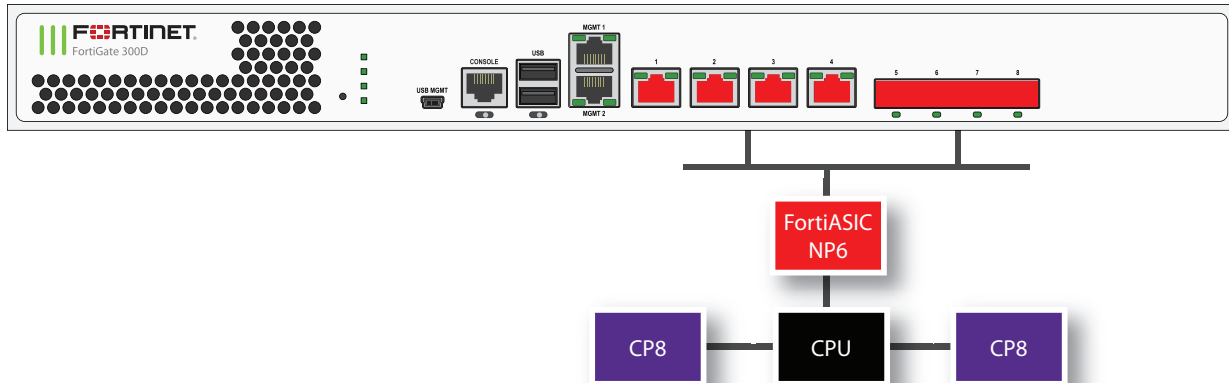
```
config system npu
  set iph-rsvd-re-cksum enable
end
```

FortiGate NP6 architectures

This chapter shows the NP6 architecture for the all FortiGate models that include NP6 processors.

FortiGate-300D fast path architecture

The FortiGate-300D includes one NP6 processor connected to four 1Gb RJ-45 Ethernet ports (port1-4) and four 1Gb SFP interfaces (port5-port8).



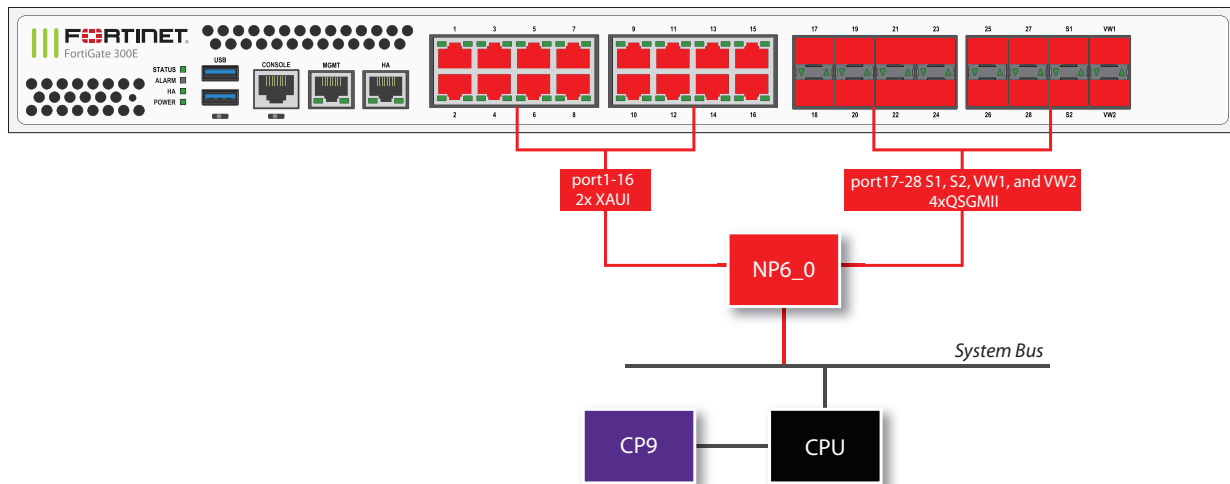
You can use the following get command to display the FortiGate-300D NP6 configuration. The command output shows one NP6 named NP6_0 and the interfaces (ports) connected to it. You can also use the `diagnose npu np6 port-list` command to display this information.

```
get hardware npu np6 port-list
Chip  XAUI Ports  Max  Cross-chip
      Speed offloading
-----
np6_0  0
      1  port5   1G   Yes
      1  port7   1G   Yes
      1  port8   1G   Yes
      1  port6   1G   Yes
      1  port3   1G   Yes
      1  port4   1G   Yes
      1  port1   1G   Yes
      1  port2   1G   Yes
      2
      3
-----
```

FortiGate-300E and 301E fast path architecture

The FortiGate-300E and 301E each include one NP6 processor. All front panel data interfaces (port1 to 28, S1, S2, VW1, and VW2) connect to the NP6 processor. So all supported traffic passing between any two data interfaces can be offloaded.

The following diagram also shows the XAUI and QSGMII port connections between the NP6 processor and the front panel interfaces.



You can use the following `get` command to display the FortiGate-300E or 301E NP6 configuration. You can also use the `diagnose npu np6 port-list` command to display this information.


```
get hardware npu np6 port-list
```

Chip	XAUI	Ports	Max Speed	Cross-chip offloading
np6_0	0	port1	1G	Yes
	0	port2	1G	Yes
	0	port3	1G	Yes
	0	port4	1G	Yes
	0	port5	1G	Yes
	0	port6	1G	Yes
	0	port7	1G	Yes
	0	port8	1G	Yes
	1	port9	1G	Yes
	1	port10	1G	Yes
	1	port11	1G	Yes
	1	port12	1G	Yes
	1	port13	1G	Yes
	1	port14	1G	Yes
	1	port15	1G	Yes
	1	port16	1G	Yes
	2	port17	1G	Yes
	2	port18	1G	Yes
	2	port19	1G	Yes
	2	port20	1G	Yes
	2	port21	1G	Yes
	2	port22	1G	Yes
	2	port23	1G	Yes
	2	port24	1G	Yes
	3	port25	1G	Yes
	3	port26	1G	Yes
	3	port27	1G	Yes
	3	port28	1G	Yes
	3	s1	1G	Yes
	3	s2	1G	Yes

3	vw1	1G	Yes
3	vw2	1G	Yes

FortiGate-400D fast path architecture

The FortiGate-400D includes one NP6 processor connected to eight 1Gb SFP interfaces (port1-port8) and eight 1Gb RJ-45 Ethernet ports (port9-16).

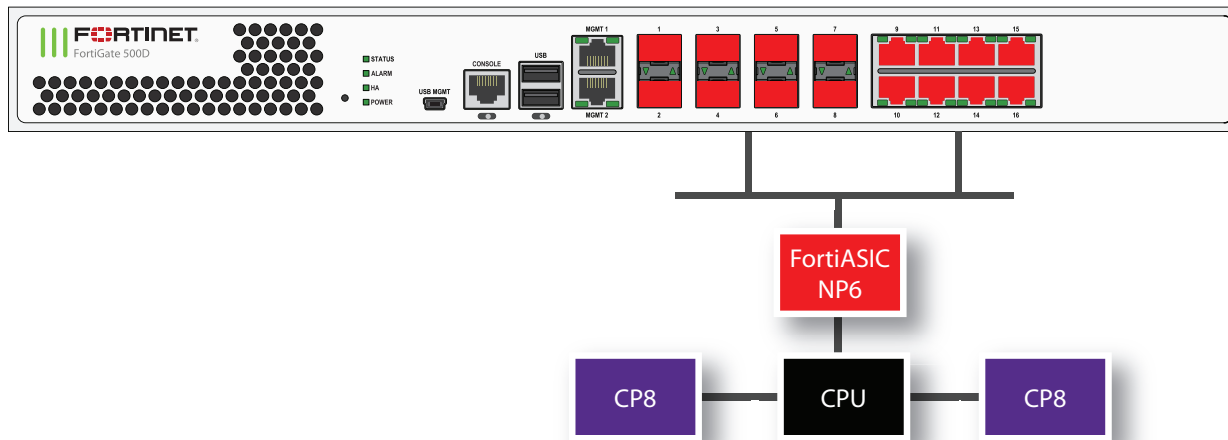


You can use the following get command to display the FortiGate-400D NP6 configuration. The command output shows one NP6 named NP6_0 and the interfaces (ports) connected to it. You can also use the `diagnose npu np6 port-list` command to display this information.

```
get hardware npu np6 port-list
Chip  XAUI Ports  Max  Cross-chip
      Speed offloading
-----
np6_0  0
      1  port10  1G    Yes
      1  port9   1G    Yes
      1  port12  1G    Yes
      1  port11  1G    Yes
      1  port14  1G    Yes
      1  port13  1G    Yes
      1  port16  1G    Yes
      1  port15  1G    Yes
      1  port5   1G    Yes
      1  port7   1G    Yes
      1  port8   1G    Yes
      1  port6   1G    Yes
      1  port3   1G    Yes
      1  port4   1G    Yes
      1  port1   1G    Yes
      1  port2   1G    Yes
      2
      3
-----
```

FortiGate-500D fast path architecture

The FortiGate-500D includes one NP6 processor connected to eight 1Gb SFP interfaces (port1-port8) and eight 1Gb RJ-45 Ethernet ports (port9-16).



You can use the following `get` command to display the FortiGate-500D NP6 configuration. The command output shows one NP6 named NP6_0 and the interfaces (ports) connected to it. You can also use the `diagnose npu np6 port-list` command to display this information.

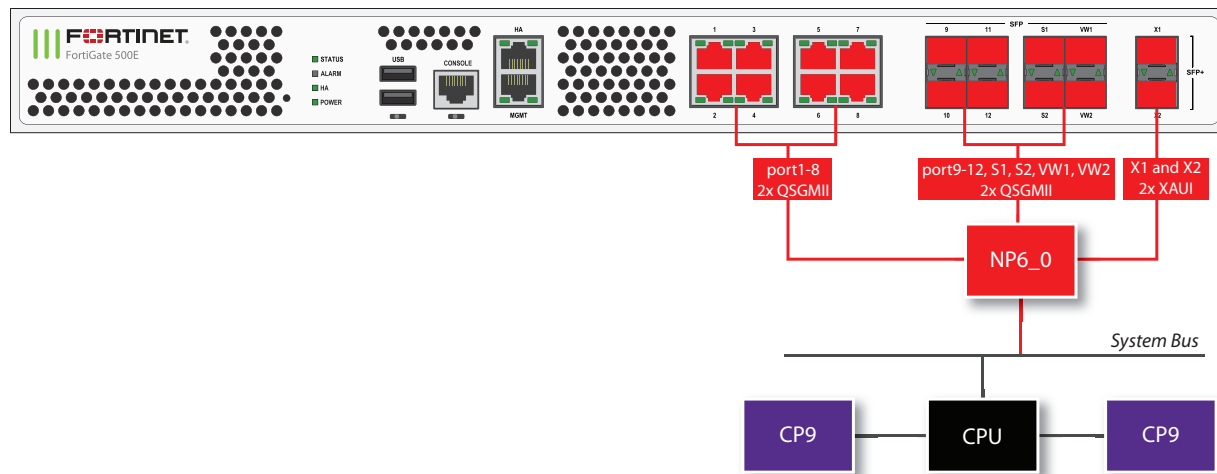
```
get hardware npu np6 port-list
Chip   XAUI Ports   Max   Cross-chip
        Speed offloading
-----
np6_0  0
      1   port10  1G    Yes
      1   port9   1G    Yes
      1   port12  1G    Yes
      1   port11  1G    Yes
      1   port14  1G    Yes
      1   port13  1G    Yes
      1   port16  1G    Yes
      1   port15  1G    Yes
      1   port5   1G    Yes
      1   port7   1G    Yes
      1   port8   1G    Yes
      1   port6   1G    Yes
      1   port3   1G    Yes
      1   port4   1G    Yes
      1   port1   1G    Yes
      1   port2   1G    Yes
      2
      3
-----
```

FortiGate-500E and 501E fast path architecture

The FortiGate-500E and 501E each include one NP6 processor. All front panel data interfaces (port1 to 212, S1, S2, VW1, VW2, X1 and X2) connect to the NP6 processor. So all supported traffic passing between any two data

interfaces can be offloaded.

The following diagram also shows the QSGMII and XAUI port connections between the NP6 processor and the front panel interfaces.




You can use the following get command to display the FortiGate-500E or 501E NP6 configuration. You can also use the diagnose npu np6 port-list command to display this information.

```
get hardware npu np6 port-list
```

Chip	XAUI	Ports	Max Speed	Cross-chip offloading
np6_0	0	x1	10G	Yes
	1	port1	1G	Yes
	1	port2	1G	Yes
	1	port3	1G	Yes
	1	port4	1G	Yes
	1	port5	1G	Yes
	1	port6	1G	Yes
	1	port7	1G	Yes
	1	port8	1G	Yes
	1	port9	1G	Yes
	1	port10	1G	Yes
	1	port11	1G	Yes
	1	port12	1G	Yes
	1	s1	1G	Yes
	1	s2	1G	Yes
	1	vw1	1G	Yes
	1	vw2	1G	Yes
	2	x2	10G	Yes
	3			

FortiGate-600D fast path architecture

The FortiGate-600D includes one NP6 processor connected to eight 1Gb SFP interfaces (port1-port8) and eight 1Gb RJ-45 Ethernet ports (port9-16) and two 10Gb SFP+ interfaces (port17 and port18).

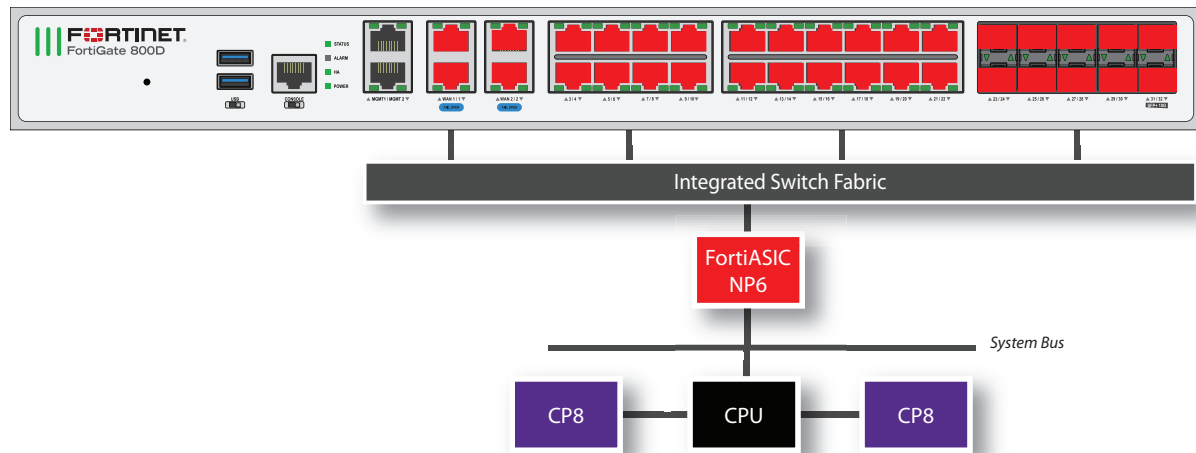


You can use the following `get` command to display the FortiGate-600D NP6 configuration. The command output shows one NP6 named NP6_0 and the interfaces (ports) connected to it. You can also use the `diagnose npu np6 port-list` command to display this information.

```
get hardware npu np6 port-list
Chip  XAUI Ports  Max  Cross-chip
      Speed offloading
-----
np6_0  0
      1  port10  1G   Yes
      1  port9   1G   Yes
      1  port12  1G   Yes
      1  port11  1G   Yes
      1  port14  1G   Yes
      1  port13  1G   Yes
      1  port16  1G   Yes
      1  port15  1G   Yes
      1  port5   1G   Yes
      1  port7   1G   Yes
      1  port8   1G   Yes
      1  port6   1G   Yes
      1  port3   1G   Yes
      1  port4   1G   Yes
      1  port1   1G   Yes
      1  port2   1G   Yes
      2  port17  10G  Yes
      3  port18  10G  Yes
-----
```

FortiGate-800D fast path architecture

The FortiGate-800D includes one NP6 processor connected through an integrated switch fabric to all of the FortiGate-800D network interfaces. This hardware configuration supports NP6-accelerated fast path offloading for sessions between any of the FortiGate-800D interfaces.



You can use the following `get` command to display the FortiGate-800D NP6 configuration. The command output shows one NP6 named NP6_0. The output also shows all of the FortiGate-800D interfaces (ports) connected to NP6_0. You can also use the `diagnose npu np6 port-list` command to display this information.

```
get hardware npu np6 port-list
```

Chip	XAUI	Ports	Max Speed	Cross-chip offloading

np6_0	0	port31	10G	Yes
	1	wan1	1G	Yes
	1	port1	1G	Yes
	1	wan2	1G	Yes
	1	port2	1G	Yes
	1	port3	1G	Yes
	1	port4	1G	Yes
	1	port5	1G	Yes
	1	port6	1G	Yes
	1	port30	1G	Yes
	1	port29	1G	Yes
	1	port28	1G	Yes
	1	port27	1G	Yes
	1	port26	1G	Yes
	1	port25	1G	Yes
	1	port24	1G	Yes
	1	port23	1G	Yes
	2	port7	1G	Yes
	2	port8	1G	Yes
	2	port9	1G	Yes
	2	port10	1G	Yes
	2	port11	1G	Yes
	2	port12	1G	Yes
	2	port13	1G	Yes
	2	port14	1G	Yes
	2	port15	1G	Yes
	2	port16	1G	Yes
	2	port17	1G	Yes
	2	port18	1G	Yes
	2	port19	1G	Yes
	2	port20	1G	Yes
	2	port21	1G	Yes
	2	port22	1G	Yes
	3	port32	10G	Yes

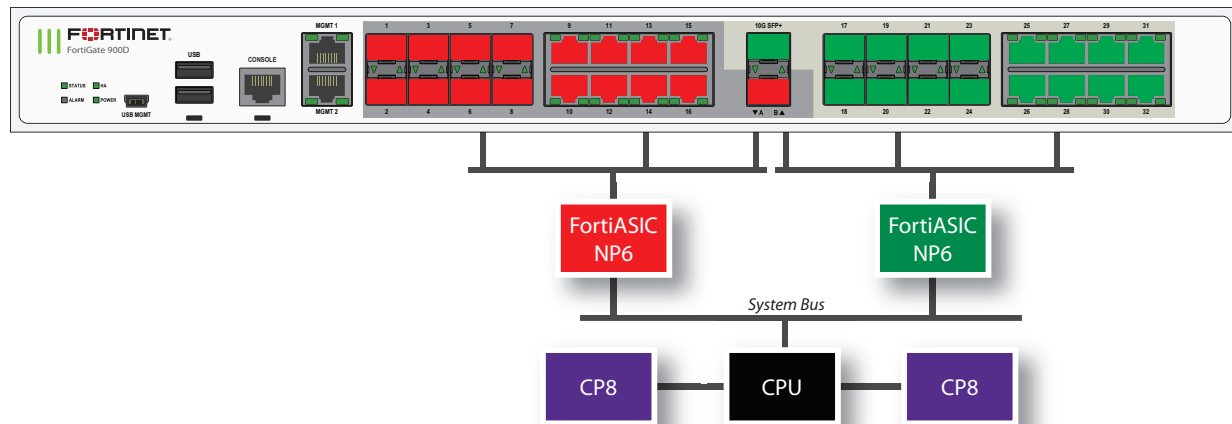
FortiGate-900D fast path architecture

The FortiGate-900D includes two NP6 processors that are not connected by an integrated switch fabric (ISF). Without an ISF, traffic through a FortiGate-900D could experience lower latency than traffic through similar hardware with an ISF. The NP6 processors are connected to network interfaces as follows:



Because the FortiGate-900D does not have an ISF you cannot create Link Aggregation Groups (LAGs) that include interfaces connected to both NP6 processors.

- Eight 1Gb SFP interfaces (port17-port24), eight 1Gb RJ-45 Ethernet interfaces (port25-32) and one 10Gb SFP+ interface (portB) share connections to the first NP6 processor.
- Eight 1Gb SFP interfaces (port1-port8), eight RJ-45 Ethernet interfaces (port9-16) and one 10Gb SFP+ interface (portA) share connections to the second NP6 processor.



You can use the following get command to display the FortiGate-900D NP6 configuration. The command output shows two NP6s named NP6_0 and NP6_1. The output also shows the interfaces (ports) connected to each NP6. You can also use the `diagnose npu np6 port-list` command to display this information.

```

get hardware npu np6 port-list
Chip  XAUI Ports  Max  Cross-chip
      Speed offloading
-----
np6_0  0
      1  port17  1G   Yes
      1  port18  1G   Yes
      1  port19  1G   Yes
      1  port20  1G   Yes
      1  port21  1G   Yes
      1  port22  1G   Yes
      1  port23  1G   Yes
      1  port24  1G   Yes
      1  port27  1G   Yes
      1  port28  1G   Yes
      1  port25  1G   Yes
      1  port26  1G   Yes
      1  port31  1G   Yes
      1  port32  1G   Yes
      1  port29  1G   Yes
      1  port30  1G   Yes
      2  portB   10G  Yes
      3
-----
np6_1  0
      1  port1   1G   Yes
      1  port2   1G   Yes
      1  port3   1G   Yes
      1  port4   1G   Yes
      1  port5   1G   Yes
      1  port6   1G   Yes
      1  port7   1G   Yes
      1  port8   1G   Yes
      1  port11  1G   Yes
      1  port12  1G   Yes
      1  port9   1G   Yes
      1  port10  1G   Yes
      1  port15  1G   Yes
      1  port16  1G   Yes
      1  port13  1G   Yes
      1  port14  1G   Yes
      2  portA   10G  Yes
      3

```


FortiGate-1000D fast path architecture

The FortiGate-1000D includes two NP6 processors that are not connected by an integrated switch fabric (ISF). The NP6 processors are connected to network interfaces as follows:



Because the FortiGate-1000D does not have an ISF you cannot create Link Aggregation Groups (LAGs) that include interfaces connected to both NP6 processors.

- Eight 1Gb SFP interfaces (port17-port24), eight 1Gb RJ-45 Ethernet interfaces (port25-32) and one 10Gb SFP+ interface (portB) share connections to the first NP6 processor.
- Eight 1Gb SFP interfaces (port1-port8), eight RJ-45 Ethernet interfaces (port9-16) and one 10Gb SFP+ interface (portA) share connections to the second NP6 processor.



You can use the following `get` command to display the FortiGate-1000D NP6 configuration. The command output shows two NP6s named NP6_0 and NP6_1. The output also shows the interfaces (ports) connected to each NP6. You can also use the `diagnose npu np6 port-list` command to display this information.

```


get hardware npu np6 port-list
Chip  XAUI Ports  Max  Cross-chip
      Speed offloading
-----
np6_0  0
      1  port17  1G  Yes
      1  port18  1G  Yes
      1  port19  1G  Yes
      1  port20  1G  Yes
      1  port21  1G  Yes
      1  port22  1G  Yes
      1  port23  1G  Yes
      1  port24  1G  Yes
      1  port27  1G  Yes
      1  port28  1G  Yes
      1  port25  1G  Yes
      1  port26  1G  Yes
      1  port31  1G  Yes
      1  port32  1G  Yes
      1  port29  1G  Yes
      1  port30  1G  Yes
      2  portB   10G  Yes
      3
-----
np6_1  0
      1  port1   1G  Yes
      1  port2   1G  Yes
      1  port3   1G  Yes
      1  port4   1G  Yes
      1  port5   1G  Yes
      1  port6   1G  Yes
      1  port7   1G  Yes
      1  port8   1G  Yes
      1  port11  1G  Yes
      1  port12  1G  Yes
      1  port9   1G  Yes
      1  port10  1G  Yes
      1  port15  1G  Yes
      1  port16  1G  Yes
      1  port13  1G  Yes
      1  port14  1G  Yes
      2  portA   10G  Yes
      3

```

FortiGate-1200D fast path architecture

The FortiGate-1200D features two NP6 processors both connected to an integrated switch fabric.

- Eight SFP 1Gb interfaces (port1-port8), eight RJ-45 Ethernet ports (port17-24) and two SFP+ 10Gb interfaces (port33 and port34) share connections to the first NP6 processor.
- Eight SFP 1Gb interfaces (port9-port16), eight RJ-45 Ethernet ports (port25-32) and two SFP+ 10Gb interfaces (port35-port36) share connections to the second NP6 processor.



You can use the following `get` command to display the FortiGate-1200D NP6 configuration. The command output shows two NP6s named NP6_0 and NP6_1. The output also shows the interfaces (ports) connected to each NP6. You can also use the `diagnose npu np6 port-list` command to display this information.

```


get hardware npu np6 port-list
Chip    XAUI Ports    Max    Cross-chip
        Speed offloading
-----
np6_0   0    port33  10G    Yes
        1    port34  10G    Yes
        2    port1   1G     Yes
        2    port3   1G     Yes
        2    port5   1G     Yes
        2    port7   1G     Yes
        2    port17  1G     Yes
        2    port19  1G     Yes
        2    port21  1G     Yes
        2    port23  1G     Yes
        3    port2   1G     Yes
        3    port4   1G     Yes
        3    port6   1G     Yes
        3    port8   1G     Yes
        3    port18  1G     Yes
        3    port20  1G     Yes
        3    port22  1G     Yes
        3    port24  1G     Yes
-----
np6_1   0    port35  10G    Yes
        1    port36  10G    Yes
        2    port9   1G     Yes
        2    port11  1G     Yes
        2    port13  1G     Yes
        2    port15  1G     Yes
        2    port25  1G     Yes
        2    port27  1G     Yes
        2    port29  1G     Yes
        2    port31  1G     Yes
        3    port10  1G     Yes
        3    port12  1G     Yes
        3    port14  1G     Yes
        3    port16  1G     Yes
        3    port26  1G     Yes
        3    port28  1G     Yes
        3    port30  1G     Yes
        3    port32  1G     Yes
-----

```


FortiGate-1500D fast path architecture

The FortiGate-1500D features two NP6 processors both connected to an integrated switch fabric.

- Eight SFP 1Gb interfaces (port1-port8), eight RJ-45 1Gb Ethernet interfaces (port17-24) and four SFP+ 10Gb interfaces (port33-port36) share connections to the first NP6 processor.
- Eight SFP 1Gb interfaces (port9-port16), eight RJ-45 1Gb Ethernet interfaces (port25-32) and four SFP+ 10Gb interfaces (port37-port40) share connections to the second NP6 processor.



You can use the following get command to display the FortiGate-1500D NP6 configuration. The command output shows two NP6s named NP6_0 and NP6_1. The output also shows the interfaces (ports) connected to each NP6. You can also use the `diagnose npu np6 port-list` command to display this information.

```
get hardware npu np6 port-list
```

Chip	XAUI	Ports	Max Speed	Cross-chip offloading
np6_0	0	port1	1G	Yes
	0	port5	1G	Yes
	0	port17	1G	Yes
	0	port21	1G	Yes
	0	port33	10G	Yes
	1	port2	1G	Yes
	1	port6	1G	Yes
	1	port18	1G	Yes
	1	port22	1G	Yes
	1	port34	10G	Yes
	2	port3	1G	Yes

	2	port7	1G	Yes
	2	port19	1G	Yes
	2	port23	1G	Yes
	2	port35	10G	Yes
	3	port4	1G	Yes
	3	port8	1G	Yes
	3	port20	1G	Yes
	3	port24	1G	Yes
	3	port36	10G	Yes


np6_1	0	port9	1G	Yes
	0	port13	1G	Yes
	0	port25	1G	Yes
	0	port29	1G	Yes
	0	port37	10G	Yes
	1	port10	1G	Yes
	1	port14	1G	Yes
	1	port26	1G	Yes
	1	port30	1G	Yes
	1	port38	10G	Yes
	2	port11	1G	Yes
	2	port15	1G	Yes
	2	port27	1G	Yes
	2	port31	1G	Yes
	2	port39	10G	Yes
	3	port12	1G	Yes
	3	port16	1G	Yes
	3	port28	1G	Yes
	3	port32	1G	Yes
	3	port40	10G	Yes

FortiGate-1500DT fast path architecture

The FortiGate-1500DT features two NP6 processors both connected to an integrated switch fabric. The FortiGate-1500DT has the same hardware configuration as the FortiGate-1500D, but with the addition of newer CPUs.

The FortiGate-1500DT includes the following interfaces and NP6 processors:

- Eight SFP 1Gb interfaces (port1-port8), eight RJ-45 1Gb Ethernet interfaces (port17-24) and four RJ-45 10Gb Ethernet interfaces (port33-port36) share connections to the first NP6 processor.
- Eight SFP 1Gb interfaces (port9-port16), eight RJ-45 1Gb Ethernet interfaces (port25-32) and four SFP+ 10Gb interfaces (port37-port40) share connections to the second NP6 processor.



You can use the following get command to display the FortiGate-1500DT NP6 configuration. The command output shows two NP6s named NP6_0 and NP6_1. The output also shows the interfaces (ports) connected to each NP6. You can also use the `diagnose npu np6 port-list` command to display this information.

```
get hardware npu np6 port-list
```

Chip	XAUI	Ports	Max Speed	Cross-chip offloading
np6_0	0	port1	1G	Yes
	0	port5	1G	Yes
	0	port17	1G	Yes
	0	port21	1G	Yes
	0	port33	10G	Yes
	1	port2	1G	Yes
	1	port6	1G	Yes
	1	port18	1G	Yes
	1	port22	1G	Yes
	1	port34	10G	Yes
	2	port3	1G	Yes
	2	port7	1G	Yes
	2	port19	1G	Yes
	2	port23	1G	Yes
	2	port35	10G	Yes
	3	port4	1G	Yes
	3	port8	1G	Yes
	3	port20	1G	Yes
	3	port24	1G	Yes
	3	port36	10G	Yes

np6_1	0	port9	1G	Yes
	0	port13	1G	Yes
	0	port25	1G	Yes
	0	port29	1G	Yes
	0	port37	10G	Yes
	1	port10	1G	Yes
	1	port14	1G	Yes
	1	port26	1G	Yes
	1	port30	1G	Yes
	1	port38	10G	Yes
	2	port11	1G	Yes
	2	port15	1G	Yes
	2	port27	1G	Yes
	2	port31	1G	Yes
	2	port39	10G	Yes
	3	port12	1G	Yes
	3	port16	1G	Yes
	3	port28	1G	Yes
	3	port32	1G	Yes
	3	port40	10G	Yes

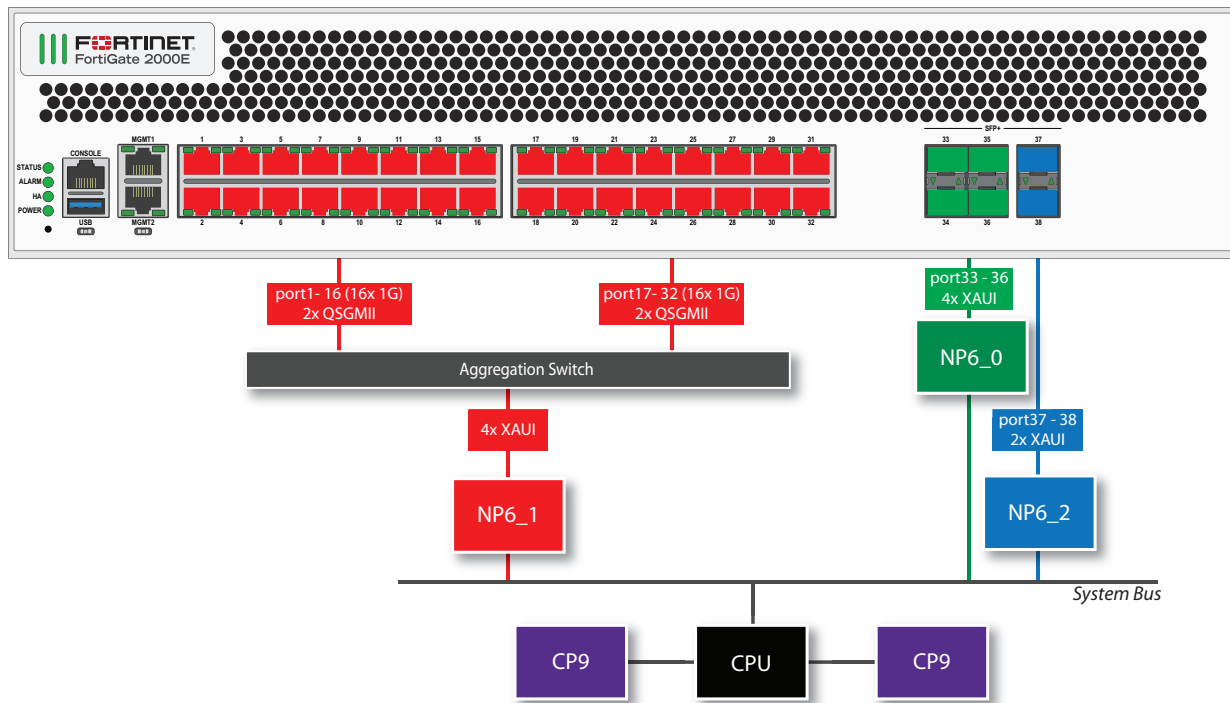
FortiGate-2000E fast path architecture

The FortiGate-2000E includes three NP6 processors in an NP Direct configuration. The NP6 processors connected to the 10GigE ports are also in a low latency NP Direct configuration. Because of NP Direct, you cannot create Link Aggregation Groups (LAGs) between interfaces connected to different NP6s. As well, traffic will only be offloaded if it enters and exits the FortiGate on interfaces connected to the same NP6.

The NP6s are connected to network interfaces as follows:

- NP6_0 is connected to four 10GigE SFP+ interfaces (port33 to port36) in a low latency configuration.
- NP6_1 is connected to thirty-two 10/100/1000BASE-T interfaces (port1 to port32).
- NP6_2 is connected to two 10GigE SFP+ (port37 and port38) in a low latency configuration.

The following diagram also shows the XAUI and QSGMII port connections between the NP6 processors and the front panel interfaces and the aggregate switch for the thirty-two 10/100/1000BASE-T interfaces.



You can use the following get command to display the FortiGate-2000E NP6 configuration. You can also use the `diagnose npu np6 port-list` command to display this information.

```

get hardware npu np6 port-list
Chip      XAUI  Ports      Max      Cross-chip
          XAUI  Ports      Speed    offloading
-----
np6_1     0      port1      1G       No
          0      port5      1G       No
          0      port9      1G       No
          0      port13     1G       No
          0      port17     1G       No
          0      port21     1G       No
          0      port25     1G       No
          0      port29     1G       No
          1      port2      1G       No
          1      port6      1G       No
          1      port10     1G       No
          1      port14     1G       No
          1      port18     1G       No
          1      port22     1G       No
          1      port26     1G       No
          1      port30     1G       No
          2      port3      1G       No
          2      port7      1G       No
          2      port11     1G       No
          2      port15     1G       No
          2      port19     1G       No
          2      port23     1G       No
          2      port27     1G       No
          2      port31     1G       No
          3      port4      1G       No
          3      port8      1G       No
          3      port12     1G       No
          3      port16     1G       No
          3      port20     1G       No
          3      port24     1G       No
          3      port28     1G       No
          3      port32     1G       No
-----
np6_0     0      port33     10G      No
          1      port34     10G      No
          2      port35     10G      No
          3      port36     10G      No
-----
np6_2     0      port37     10G      No
          1      port38     10G      No
-----

```

FortiGate-2500E fast path architecture

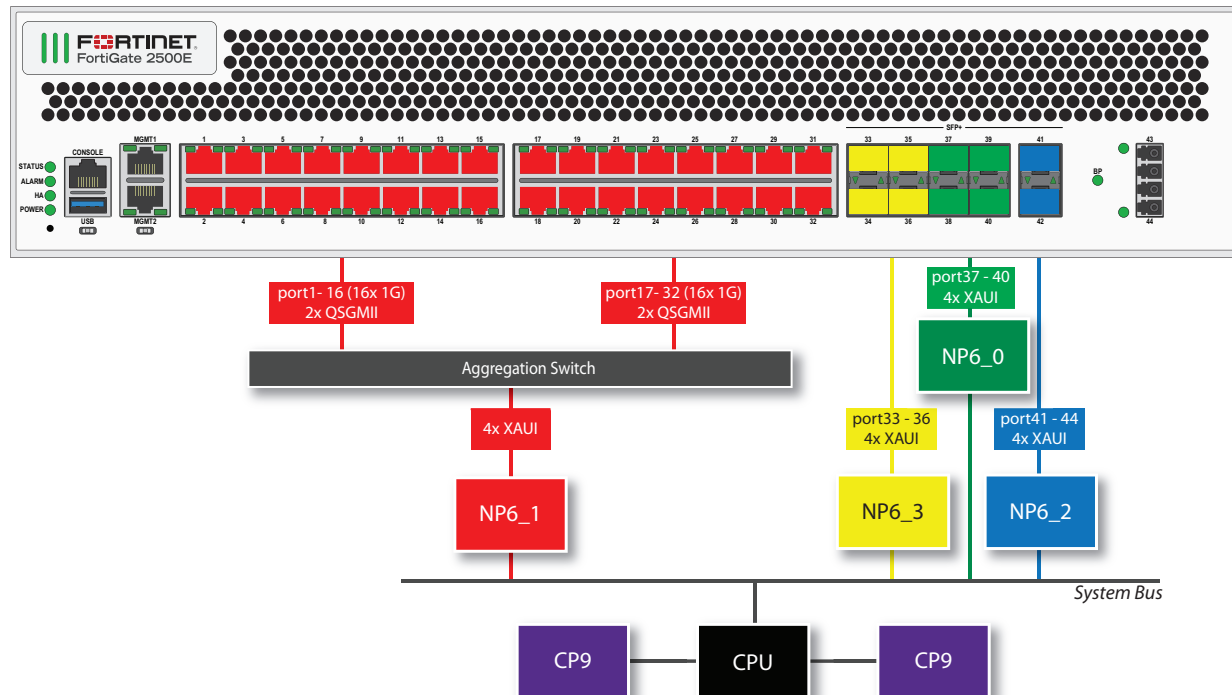
The FortiGate-2500E includes four NP6 processors in an NP Direct configuration. The NP6 processors connected to the 10GigE ports are also in a low latency NP Direct configuration. Because of NP Direct, you cannot create

Link Aggregation Groups (LAGs) between interfaces connected to different NP6s. As well, traffic will only be offloaded if it enters and exits the FortiGate on interfaces connected to the same NP6.

The NP6s are connected to network interfaces as follows:

- NP6_0 is connected to four 10GigE SFP+ interfaces (port37 to port40) in a low latency configuration.
- NP6_1 is connected to thirty-two 10/100/1000BASE-T interfaces (port1 to port32).
- NP6_2 is connected to two 10GigE SFP+ interfaces (port41 and port42) and two 10 Gig fiber bypass interfaces (port43 and port44) in a low latency configuration.
- NP6_3 is connected to four 10GigE SFP+ interfaces (port33 to port36) in a low latency configuration.

The following diagram also shows the XAUI and QSGMII port connections between the NP6 processors and the front panel interfaces and the aggregate switch for the thirty-two 10/100/1000BASE-T interfaces.



You can use the following get command to display the FortiGate-2500E NP6 configuration. You can also use the `diagnose npu np6 port-list` command to display this information.

```
get hardware npu np6 port-list
```

Chip	XAUI	Ports	Max Speed	Cross-chip offloading

np6_1	0	port1	1G	No
	0	port5	1G	No
	0	port9	1G	No
	0	port13	1G	No
	0	port17	1G	No
	0	port21	1G	No
	0	port25	1G	No
	0	port29	1G	No
	1	port2	1G	No
	1	port6	1G	No
	1	port10	1G	No
	1	port14	1G	No
	1	port18	1G	No
	1	port22	1G	No
	1	port26	1G	No
	1	port30	1G	No
	2	port3	1G	No
	2	port7	1G	No
	2	port11	1G	No
	2	port15	1G	No
	2	port19	1G	No
	2	port23	1G	No
	2	port27	1G	No
	2	port31	1G	No
	3	port4	1G	No
	3	port8	1G	No
	3	port12	1G	No
	3	port16	1G	No
	3	port20	1G	No
	3	port24	1G	No
	3	port28	1G	No
	3	port32	1G	No

np6_0	0	port37	10G	No
	1	port38	10G	No
	2	port39	10G	No
	3	port40	10G	No

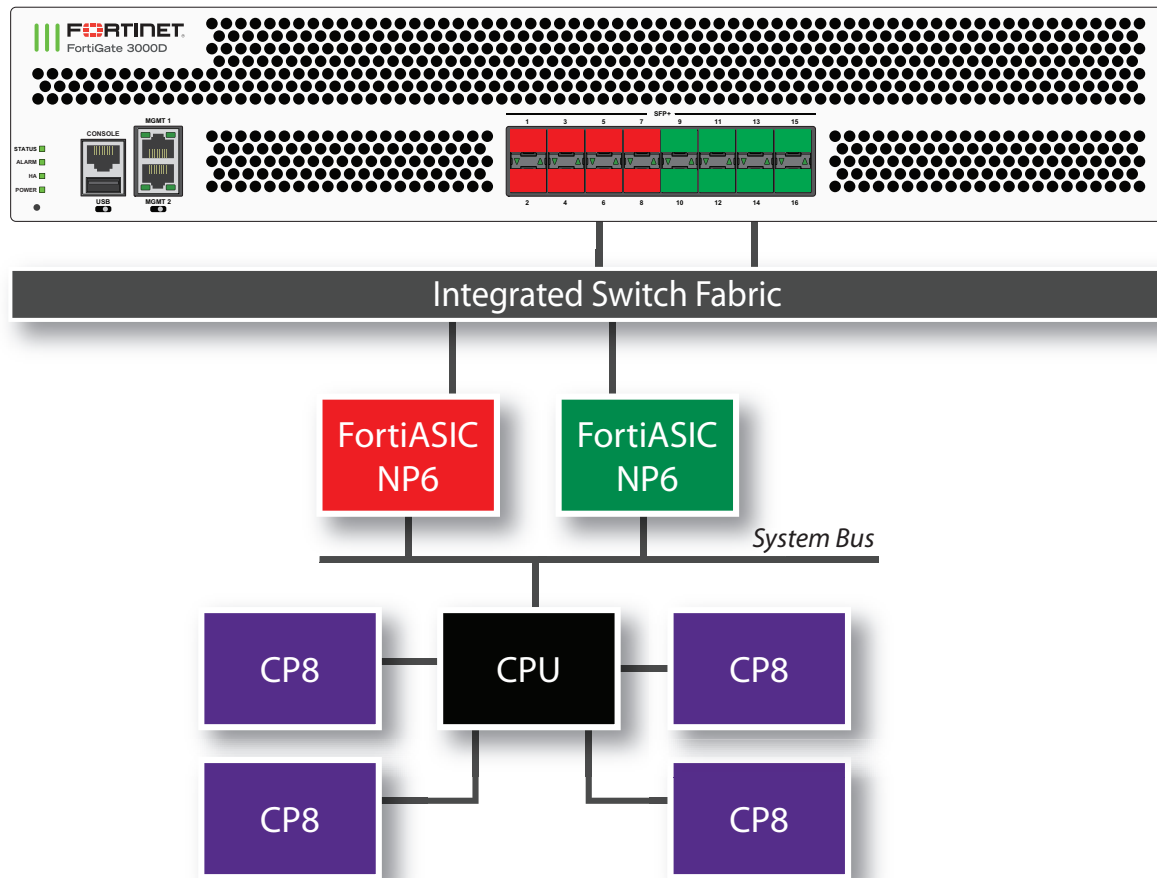
np6_2	0	port43	10G	No
	1	port44	10G	No
	2	port41	10G	No
	3	port42	10G	No

np6_3	0	port33	10G	No
	1	port34	10G	No
	2	port35	10G	No
	3	port36	10G	No

FortiGate-3000D fast path architecture

The FortiGate-3000D features 16 front panel SFP+ 10Gb interfaces connected to two NP6 processors through an Integrated Switch Fabric (ISF). The FortiGate-3000D has the following fastpath architecture:

- 8 SFP+ 10Gb interfaces, port1 through port8 share connections to the first NP6 processor (np6_0).
- 8 SFP+ 10Gb interfaces, port9 through port16 share connections to the second NP6 processor (np6_1).



You can use the following get command to display the FortiGate-3000D NP6 configuration. The command output shows two NP6s named NP6_0 and NP6_1 and the interfaces (ports) connected to each NP6. You can also use the diagnose npu np6 port-list command to display this information.

```
get hardware npu np6 port-list
```


Chip	XAUI	Ports	Max Speed	Cross-chip offloading
np6_0	0	port1	10G	Yes
	0	port6	10G	Yes
	1	port2	10G	Yes
	1	port5	10G	Yes
	2	port3	10G	Yes
	2	port8	10G	Yes
	3	port4	10G	Yes

	3	port7	10G	Yes
np6_1	0	port10	10G	Yes
	0	port13	10G	Yes
	1	port9	10G	Yes
	1	port14	10G	Yes
	2	port12	10G	Yes
	2	port15	10G	Yes
	3	port11	10G	Yes
	3	port16	10G	Yes

FortiGate-3100D fast path architecture

The FortiGate-3100D features 32 SFP+ 10Gb interfaces connected to two NP6 processors through an Integrated Switch Fabric (ISF). The FortiGate-3100D has the following fastpath architecture:

- 16 SFP+ 10Gb interfaces, port1 through port16 share connections to the first NP6 processor (np6_0).
- 16 SFP+ 10Gb interfaces, port27 through port32 share connections to the second NP6 processor (np6_1).



You can use the following get command to display the FortiGate-3100D NP6 configuration. The command output shows two NP6s named NP6_0 and NP6_1 and the interfaces (ports) connected to each NP6. You can also use the `diagnose npu np6 port-list` command to display this information.

```
get hardware npu np6 port-list
```

Chip	XAUI	Ports	Max Speed	Cross-chip offloading


np6_0	0	port1	10G	Yes
	0	port6	10G	Yes
	0	port10	10G	Yes
	0	port13	10G	Yes
	1	port2	10G	Yes
	1	port5	10G	Yes
	1	port9	10G	Yes
	1	port14	10G	Yes
	2	port3	10G	Yes
	2	port8	10G	Yes
	2	port12	10G	Yes
	2	port15	10G	Yes
	3	port4	10G	Yes
	3	port7	10G	Yes
	3	port11	10G	Yes
	3	port16	10G	Yes

np6_1	0	port17	10G	Yes
	0	port21	10G	Yes
	0	port25	10G	Yes
	0	port29	10G	Yes
	1	port18	10G	Yes
	1	port22	10G	Yes
	1	port26	10G	Yes
	1	port30	10G	Yes
	2	port19	10G	Yes
	2	port23	10G	Yes
	2	port27	10G	Yes
	2	port31	10G	Yes
	3	port20	10G	Yes
	3	port24	10G	Yes
	3	port28	10G	Yes
	3	port32	10G	Yes

FortiGate-3200D fast path architecture

The FortiGate-3200D features two NP6 processors connected to an Integrated Switch Fabric (ISF). The FortiGate-3200D has the following fastpath architecture:

- 24 SFP+ 10Gb interfaces, port1 through port24 share connections to the first NP6 processor (np6_0).
- 24 SFP+ 10Gb interfaces, port25 through port48 share connections to the second NP6 processor (np6_1).



You can use the following get command to display the FortiGate-3200D NP6 configuration. The command output shows two NP6s named NP6_0 and NP6_1 and the interfaces (ports) connected to each NP6. You can also use the `diagnose npu np6 port-list` command to display this information.

```
get hardware npu np6 port-list
```

Chip	XAUI	Ports	Max Speed	Cross-chip offloading
np6_0	0	port1	10G	Yes
	0	port5	10G	Yes
	0	port10	10G	Yes
	0	port13	10G	Yes
	0	port17	10G	Yes
	0	port22	10G	Yes
	1	port2	10G	Yes
	1	port6	10G	Yes
	1	port9	10G	Yes
	1	port14	10G	Yes
	1	port18	10G	Yes
	1	port21	10G	Yes
	2	port3	10G	Yes
	2	port7	10G	Yes
2	port12	10G	Yes	

	2	port15	10G	Yes
	2	port19	10G	Yes
	2	port24	10G	Yes
	3	port4	10G	Yes
	3	port8	10G	Yes
	3	port11	10G	Yes
	3	port16	10G	Yes
	3	port20	10G	Yes
	3	port23	10G	Yes

np6_1	0	port26	10G	Yes
	0	port29	10G	Yes
	0	port33	10G	Yes
	0	port37	10G	Yes
	0	port41	10G	Yes
	0	port45	10G	Yes
	1	port25	10G	Yes
	1	port30	10G	Yes
	1	port34	10G	Yes
	1	port38	10G	Yes
	1	port42	10G	Yes
	1	port46	10G	Yes
	2	port28	10G	Yes
	2	port31	10G	Yes
	2	port35	10G	Yes
	2	port39	10G	Yes
	2	port43	10G	Yes
	2	port47	10G	Yes
	3	port27	10G	Yes
	3	port32	10G	Yes
	3	port36	10G	Yes
	3	port40	10G	Yes
	3	port44	10G	Yes
	3	port48	10G	Yes

FortiGate-3700D fast path architecture

The FortiGate-3700D features four NP6 processors. The first two NP6 processors (np6_0 and np6_1) can be configured for low latency operation. The low latency configuration changes the FortiGate-3700D fast path architecture.

FortiGate-3700D low latency fast path architecture

Ports 25 to 32 can be used for low latency offloading. As long as traffic enters and exits the FortiGate-3700D through ports connected to the same NP6 processor and using these low latency ports the traffic will be offloaded and have lower latency than other NP6 offloaded traffic. Latency is reduced by bypassing the integrated switch fabric (ISF).

You can use the following command to turn on low latency mode for np6_0 and np6_1:

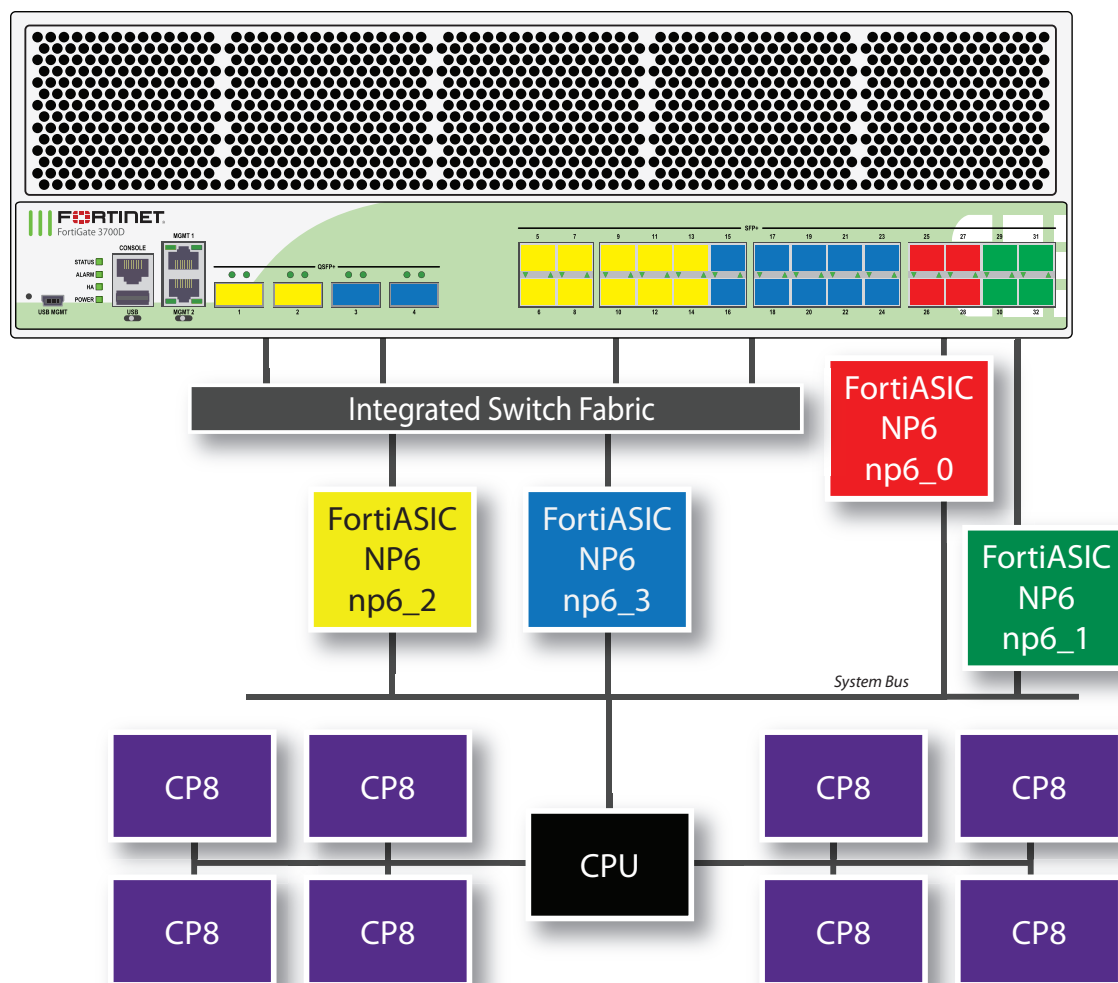
```
config system np6
  edit np6_0
    set low-latency-mode enable
  next
  edit np6_1
    set low-latency-mode enable
end
```



You do not have to turn on low latency to both np6_0 and np6_1. If you turn on low latency for just one NP6, the other NP6 will still be mapped according to the normal latency configuration.

With low latency enabled for both np6_0 and np6_1 the FortiGate-3700D has the following fastpath architecture:

- Four SFP+ 10Gb interfaces, port25 to port28, share connections to the first NP6 processor (np6_0) so sessions entering one of these ports and exiting through another will experience low latency
- Four SFP+ 10Gb interfaces, port29 to port32, share connections to the second NP6 processor (np6_1) so sessions entering one of these ports and exiting through another will experience low latency
- Ten SFP+ 10Gb interfaces, port5 to port14, and two 40Gb QSFP interfaces, port1 and port2, share connections to the third NP6 processor (np6_2).
- Ten SFP+ 10Gb interfaces, port15 to port24, and two 40Gb QSFP interfaces, port3 and port4, share connections to the fourth NP6 processor (np6_3).



You can use the following get command to display the FortiGate-3700D NP6 configuration. In this output example, the first two NP6s (np6_0 and np6_1) are configured for low latency. The command output shows four NP6s named NP6_0, NP6_1, NP6_2, and NP6_3 and the interfaces (ports) connected to each NP6. You can also use the `diagnose npu np6 port-list` command to display this information.

```
get hardware npu np6 port-list
```

Chip	XAUI Ports	Max Speed	Cross-chip offloading
np6_2	0 port5	10G	Yes
	0 port9	10G	Yes
	0 port13	10G	Yes
	1 port6	10G	Yes
	1 port10	10G	Yes
	1 port14	10G	Yes
	2 port7	10G	Yes
	2 port11	10G	Yes
	3 port8	10G	Yes
	3 port12	10G	Yes
	0-3 port1	40G	Yes

	0-3	port2	40G	Yes

np6_3	0	port15	10G	Yes
	0	port19	10G	Yes
	0	port23	10G	Yes
	1	port16	10G	Yes
	1	port20	10G	Yes
	1	port24	10G	Yes
	2	port17	10G	Yes
	2	port21	10G	Yes
	3	port18	10G	Yes
	3	port22	10G	Yes
	0-3	port3	40G	Yes
	0-3	port4	40G	Yes

np6_0	0	port26	10G	No
	1	port25	10G	No
	2	port28	10G	No
	3	port27	10G	No

np6_1	0	port30	10G	No
	1	port29	10G	No
	2	port32	10G	No
	3	port31	10G	No

FortiGate-3700D normal latency fast path architecture

You can use the following command to turn off low latency mode for np6_0 and np6_1:

```
config system np6
  edit np6_0
    set low-latency-mode disable
  next
  edit np6_1
    set low-latency-mode disable
end
```




You do not have to turn off low latency to both np6_0 and np6_1. If you turn off low latency to just one NP6, the other NP6 will still be mapped according to the normal configuration.

In addition to turning off low latency, entering these commands also changes how ports are mapped to NP6s. Port1 is now mapped to np6_0 and port 3 is not mapped to np6_1. The FortiGate-3700D has the following fastpath architecture:

- One 40Gb QSFP interface, port1, and four SFP+ 10Gb interfaces, port25 to port28 share connections to the first NP6 processor (np6_0).
- One 40Gb QSFP interface, port3, and four SFP+ 10Gb interfaces, port29 to port32 share connections to the second NP6 processor (np6_1).
- One 40Gb QSFP interface, port2 and ten SFP+ 10Gb interfaces, port5 to port14 share connections to the third NP6 processor (np6_2).

- One 40Gb QSFP interface, port4, and ten SFP+ 10Gb interfaces, port15 to port24 share connections to the fourth NP6 processor (np6_3).



You can use the following get command to display the FortiGate-3700D NP6 configuration with low latency turned off for np6_0 and np6_1. The command output shows four NP6s named NP6_0, NP6_1, NP6_2, and NP6_3 and the interfaces (ports) connected to each NP6. You can also use the `diagnose npu np6 port-list` command to display this information.

```
get hardware npu np6 port-list
```

Chip	XAUI	Ports	Max Speed	Cross-chip offloading
np6_0	0	port26	10G	Yes
	1	port25	10G	Yes
	2	port28	10G	Yes
	3	port27	10G	Yes
	0-3	port1	40G	Yes
np6_1	0	port30	10G	Yes
	1	port29	10G	Yes
	2	port32	10G	Yes

	3	port31	10G	Yes
	0-3	port3	40G	Yes

np6_2	0	port5	10G	Yes
	0	port9	10G	Yes
	0	port13	10G	Yes
	1	port6	10G	Yes
	1	port10	10G	Yes
	1	port14	10G	Yes
	2	port7	10G	Yes
	2	port11	10G	Yes
	3	port8	10G	Yes
	3	port12	10G	Yes
	0-3	port2	40G	Yes

np6_3	0	port15	10G	Yes
	0	port19	10G	Yes
	0	port23	10G	Yes
	1	port16	10G	Yes
	1	port20	10G	Yes
	1	port24	10G	Yes
	2	port17	10G	Yes
	2	port21	10G	Yes
	3	port18	10G	Yes
	3	port22	10G	Yes
	0-3	port4	40G	Yes

FortiGate-3700DX fast path architecture

The FortiGate-3700DX features four NP6 processors. The first two NP6 processors (np6_0 and np6_1) can be configured for low latency operation. The low latency configuration changes the FortiGate-3700D fast path architecture. The FortiGate-3700DX also includes two TP2 cards that offload GTPu sessions.

FortiGate-3700DX low latency fast path architecture

Ports 25 to 32 can be used for low latency offloading. As long as traffic enters and exits the FortiGate-3700D through ports connected to the same NP6 processor and using these low latency ports the traffic will be offloaded and have lower latency than other NP6 offloaded traffic. Latency is reduced by bypassing the integrated switch fabric (ISF).

You can use the following command to turn on low latency mode for np6_0 and np6_1:

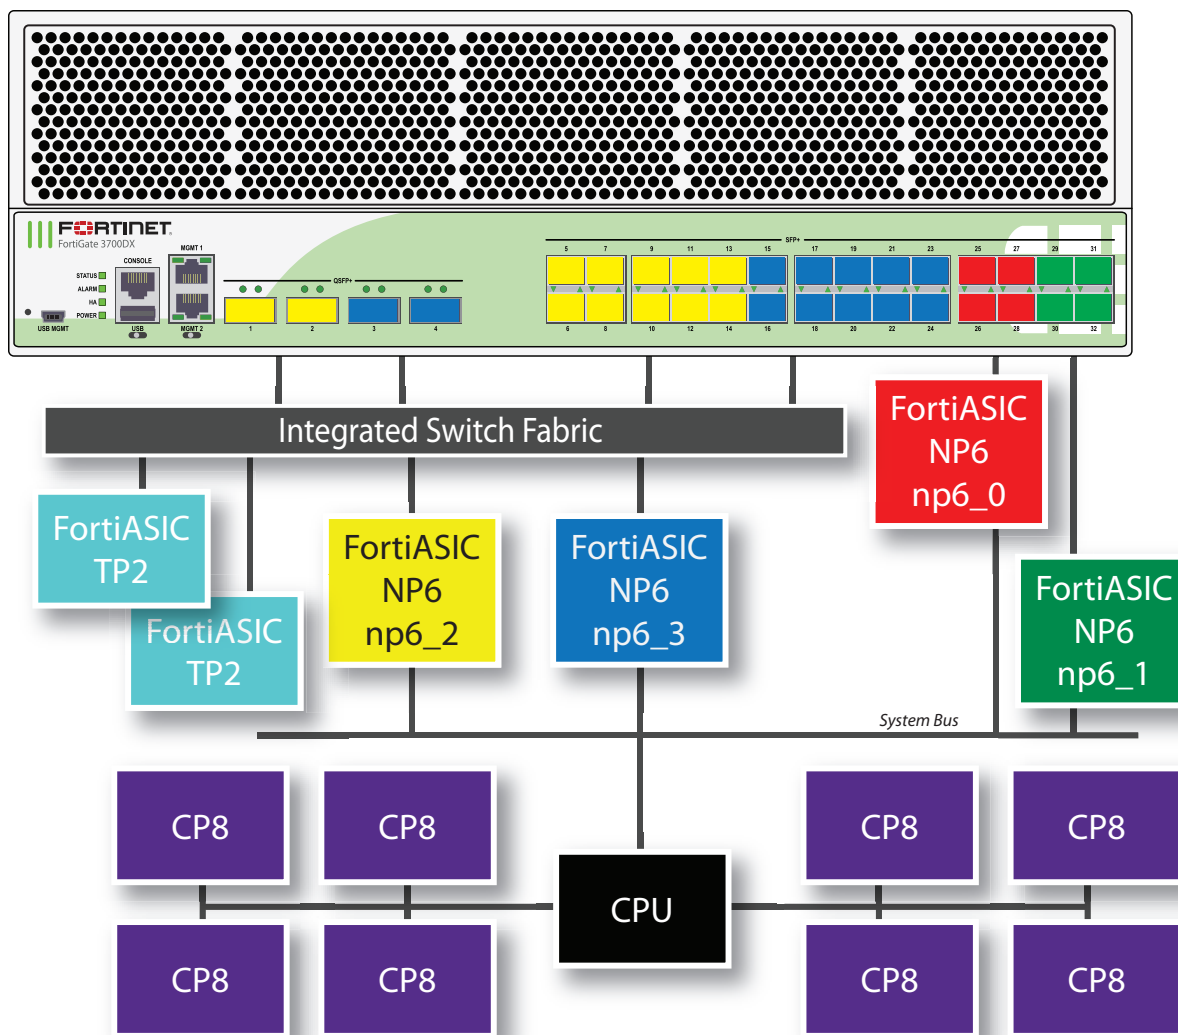
```
config system np6
  edit np6_0
    set low-latency-mode enable
  next
  edit np6_1
    set low-latency-mode enable
end
```



You do not have to turn on low latency to both np6_0 and np6_1. If you turn on low latency for just one NP6, the other NP6 will still be mapped according to the normal latency configuration.

With low latency enabled for both np6_0 and np6_1 the FortiGate-3700D has the following fastpath architecture:

- Four SFP+ 10Gb interfaces, port25 to port28, share connections to the first NP6 processor (np6_0) so sessions entering one of these ports and exiting through another will experience low latency
- Four SFP+ 10Gb interfaces, port29 to port32, share connections to the second NP6 processor (np6_1) so sessions entering one of these ports and exiting through another will experience low latency
- Ten SFP+ 10Gb interfaces, port5 to port14, and two 40Gb QSFP interfaces, port1 and port2, share connections to the third NP6 processor (np6_2).
- Ten SFP+ 10Gb interfaces, port15 to port24, and two 40Gb QSFP interfaces, port3 and port4, share connections to the fourth NP6 processor (np6_3).



You can use the following get command to display the FortiGate-3700D NP6 configuration. In this output example, the first two NP6s (np6_0 and np6_1) are configured for low latency. The command output shows four

NP6s named NP6_0, NP6_1, NP6_2, and NP6_3 and the interfaces (ports) connected to each NP6. You can also use the `diagnose npu np6 port-list` command to display this information.

```
get hardware npu np6 port-list
```

Chip	XAUI	Ports	Max Speed	Cross-chip offloading

np6_2	0	port5	10G	Yes
	0	port9	10G	Yes
	0	port13	10G	Yes
	1	port6	10G	Yes
	1	port10	10G	Yes
	1	port14	10G	Yes
	2	port7	10G	Yes
	2	port11	10G	Yes
	3	port8	10G	Yes
	3	port12	10G	Yes
	0-3	port1	40G	Yes
	0-3	port2	40G	Yes

np6_3	0	port15	10G	Yes
	0	port19	10G	Yes
	0	port23	10G	Yes
	1	port16	10G	Yes
	1	port20	10G	Yes
	1	port24	10G	Yes
	2	port17	10G	Yes
	2	port21	10G	Yes
	3	port18	10G	Yes
	3	port22	10G	Yes
	0-3	port3	40G	Yes
	0-3	port4	40G	Yes

np6_0	0	port26	10G	No
	1	port25	10G	No
	2	port28	10G	No
	3	port27	10G	No

np6_1	0	port30	10G	No
	1	port29	10G	No
	2	port32	10G	No
	3	port31	10G	No

FortiGate-3700D normal latency fast path architecture

You can use the following command to turn off low latency mode for np6_0 and np6_1:

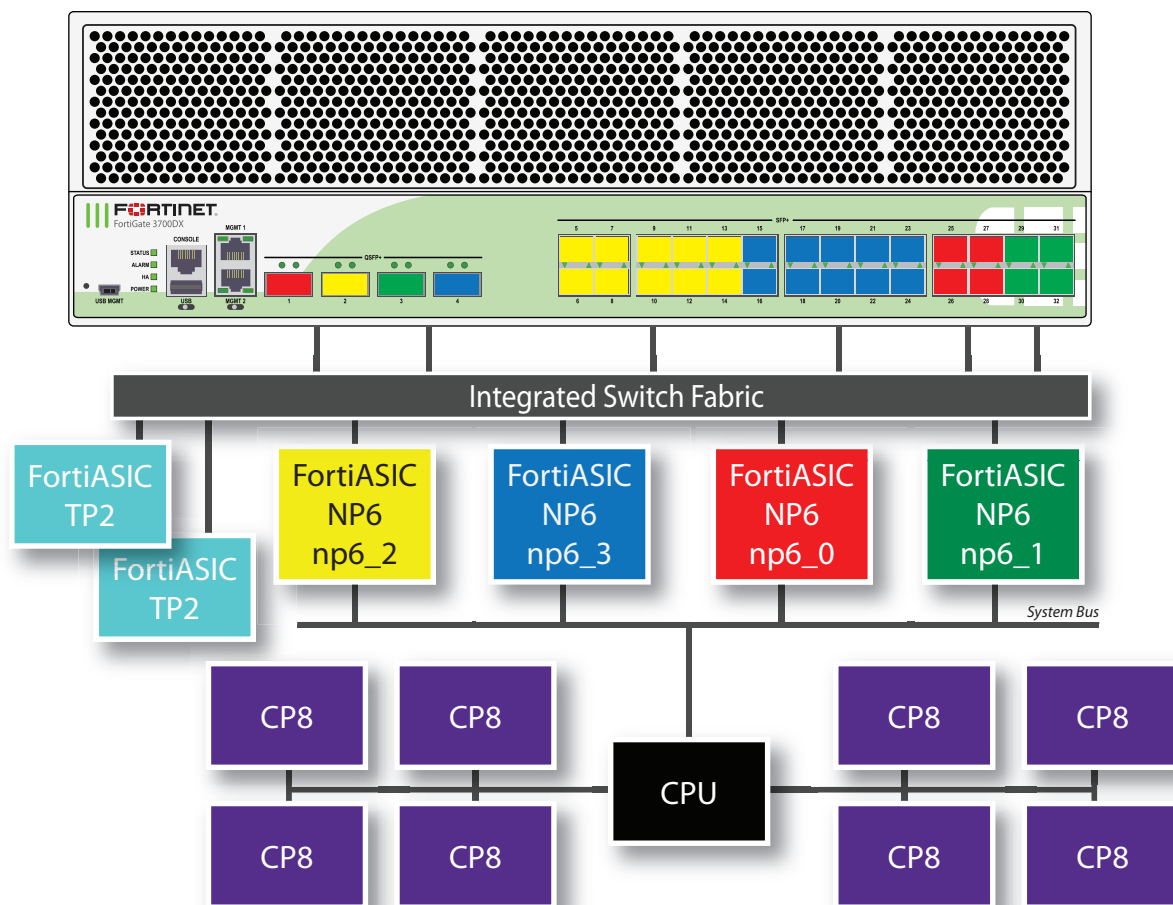
```
config system np6
  edit np6_0
    set low-latency-mode disable
  next
  edit np6_1
    set low-latency-mode disable
end
```



You do not have to turn off low latency to both np6_0 and np6_1. If you turn off low latency to just one NP6, the other NP6 will still be mapped according to the normal configuration.

In addition to turning off low latency, entering these commands also changes how ports are mapped to NP6s. Port1 is now mapped to np6_0 and port 3 is not mapped to np6_1. The FortiGate-3700D has the following fastpath architecture:

- One 40Gb QSFP interface, port1, and four SFP+ 10Gb interfaces, port25 to port28 share connections to the first NP6 processor (np6_0).
- One 40Gb QSFP interface, port3, and four SFP+ 10Gb interfaces, port29 to port32 share connections to the second NP6 processor (np6_1).
- One 40Gb QSFP interface, port2 and ten SFP+ 10Gb interfaces, port5 to port14 share connections to the third NP6 processor (np6_2).
- One 40Gb QSFP interface, port4, and ten SFP+ 10Gb interfaces, port15 to port24 share connections to the fourth NP6 processor (np6_3).



You can use the following get command to display the FortiGate-3700D NP6 configuration with low latency turned off for np6_0 and np6_1. The command output shows four NP6s named NP6_0, NP6_1, NP6_2, and NP6_3 and the interfaces (ports) connected to each NP6. You can also use the `diagnose npu np6 port-list` command to display this information.

```
get hardware npu np6 port-list
```

Chip	XAUI	Ports	Max Speed	Cross-chip offloading

np6_0	0	port26	10G	Yes
	1	port25	10G	Yes
	2	port28	10G	Yes
	3	port27	10G	Yes
	0-3	port1	40G	Yes


np6_1	0	port30	10G	Yes
	1	port29	10G	Yes
	2	port32	10G	Yes
	3	port31	10G	Yes
	0-3	port3	40G	Yes

np6_2	0	port5	10G	Yes
	0	port9	10G	Yes
	0	port13	10G	Yes
	1	port6	10G	Yes
	1	port10	10G	Yes
	1	port14	10G	Yes
	2	port7	10G	Yes
	2	port11	10G	Yes
	3	port8	10G	Yes
	3	port12	10G	Yes
	0-3	port2	40G	Yes

np6_3	0	port15	10G	Yes
	0	port19	10G	Yes
	0	port23	10G	Yes
	1	port16	10G	Yes
	1	port20	10G	Yes
	1	port24	10G	Yes
	2	port17	10G	Yes
	2	port21	10G	Yes
	3	port18	10G	Yes
	3	port22	10G	Yes
	0-3	port4	40G	Yes

FortiGate-3800D fast path architecture

The FortiGate-3800D features four front panel 100GigE CFP2 interfaces, four 40GigE QSFP+ interfaces, and eight 10GigE SFP+ interfaces connected to eight NP6 processors through an Integrated Switch Fabric (ISF). Individual interfaces are not mapped to NP6 processors because of the integrated switch fabric. No special mapping is required for fastpath offloading or aggregate interfaces.



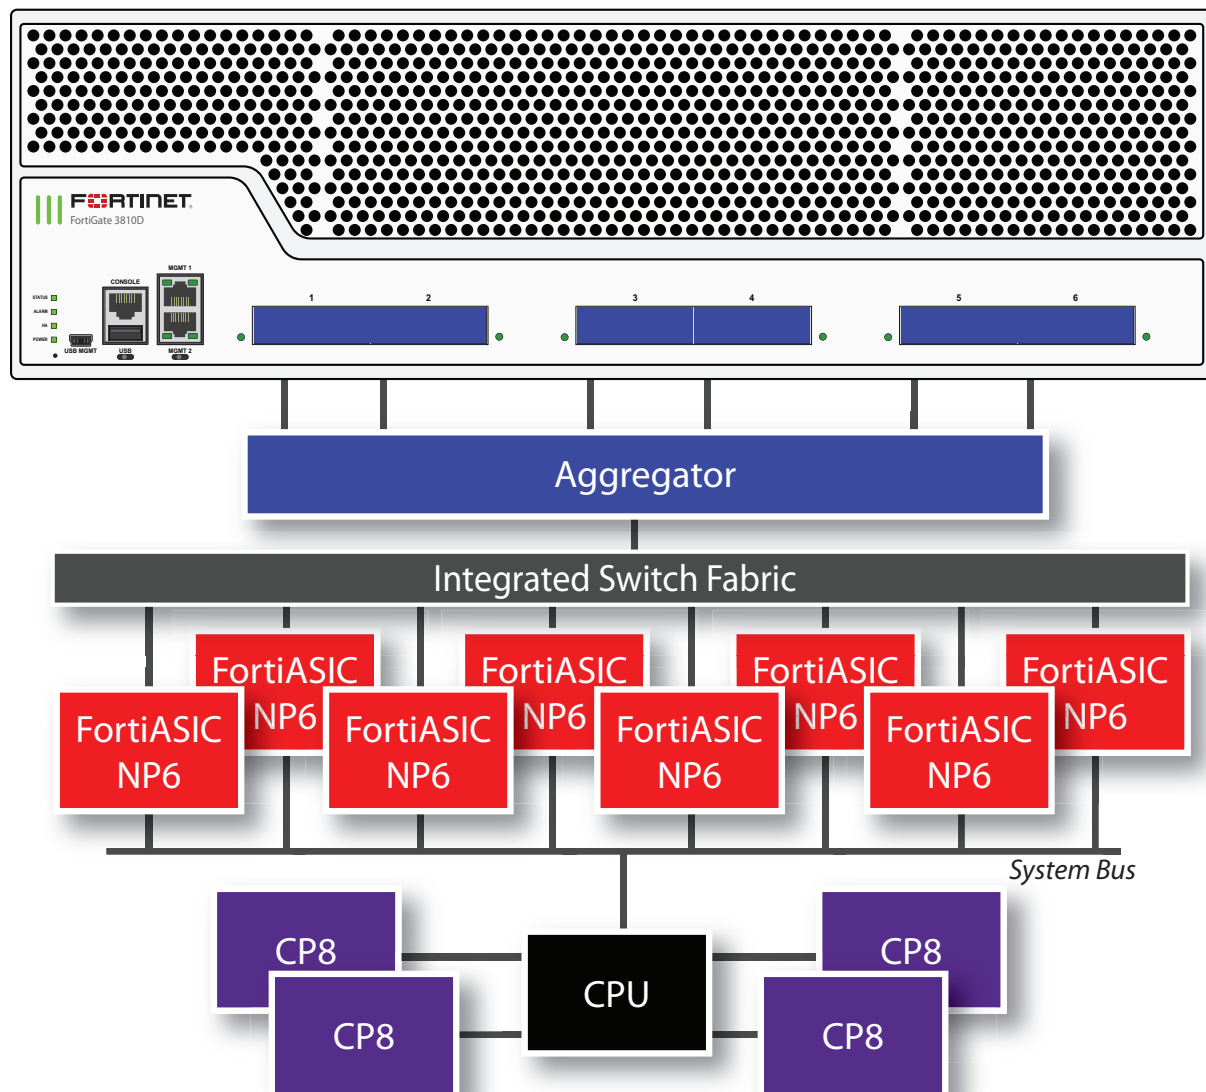
You can use the following get command to display the FortiGate-3800D NP6 configuration. The command output shows all NP6s connected to each interface (port) with cross-chip offloading supported for each port. You can also use the diagnose npu np6 port-list command to display this information.

Chip	XAUI	Ports	Max Speed	Cross-chip offloading
NP#0-7	0-3	port1	100000M	Yes
NP#0-7	0-3	port2	100000M	Yes
NP#0-7	0-3	port3	100000M	Yes
NP#0-7	0-3	port4	100000M	Yes
NP#0-7	0-3	port5	40000M	Yes
NP#0-7	0-3	port6	40000M	Yes
NP#0-7	0-3	port7	40000M	Yes
NP#0-7	0-3	port8	40000M	Yes
NP#0-7	0-3	port9	10000M	Yes
NP#0-7	0-3	port10	10000M	Yes

NP#0-7	0-3	port11	10000M	Yes
NP#0-7	0-3	port12	10000M	Yes
NP#0-7	0-3	port13	10000M	Yes
NP#0-7	0-3	port14	10000M	Yes
NP#0-7	0-3	port15	10000M	Yes
NP#0-7	0-3	port16	10000M	Yes
-----	-----	-----	-----	-----

FortiGate-3810D fast path architecture

The FortiGate-3810D features six front panel 100GigE CFP2 interfaces connected to eight NP6 processors through an Integrated Switch Fabric (ISF). Individual interfaces are not mapped to NP6 processors because of the integrated switch fabric. No special mapping is required for fastpath offloading or aggregate interfaces.



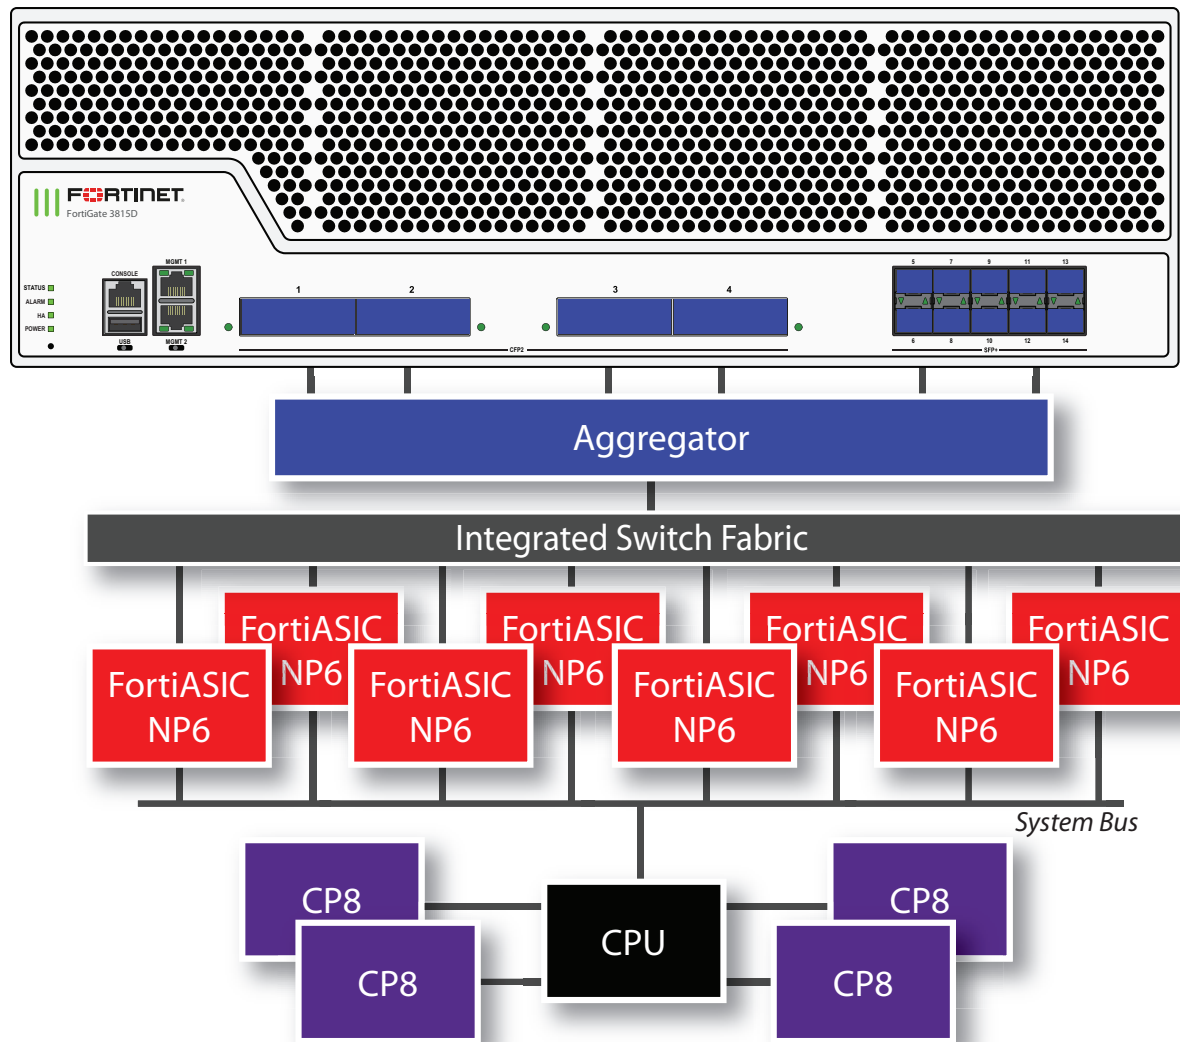
You can use the following get command to display the FortiGate-3810D NP6 configuration. The command output shows all NP6s connected to each interface (port) with cross-chip offloading supported for each port. You can also use the `diagnose npu np6 port-list` command to display this information.

```
get hardware npu np6 port-list
```

Chip	XAUI	Ports	Max Speed	Cross-chip offloading
all	0-3	port1	100000M	Yes
all	0-3	port2	100000M	Yes
all	0-3	port3	100000M	Yes
all	0-3	port4	100000M	Yes
all	0-3	port5	100000M	Yes
all	0-3	port6	100000M	Yes

FortiGate-3815D fast path architecture

The FortiGate-3815D features four front panel 100GigE CFP2 interfaces and eight 10GigE SFP+ interfaces connected to eight NP6 processors through an Integrated Switch Fabric (ISF). Individual interfaces are not mapped to NP6 processors because of the integrated switch fabric. No special mapping is required for fastpath offloading or aggregate interfaces.




You can use the following get command to display the FortiGate-3815D NP6 configuration. The command output shows all NP6s connected to each interface (port) with cross-chip offloading supported for each port. You can also use the diagnose npu np6 port-list command to display this information.

```
get hardware npu np6 port-list
Chip   XAUI Ports   Max      Cross-chip
                  Speed    offloading
-----
all    0-3   port1    100000M  Yes
all    0-3   port2    100000M  Yes
all    0-3   port3    100000M  Yes
all    0-3   port4    100000M  Yes
all    0-3   port11   10000M   Yes
all    0-3   port12   10000M   Yes
all    0-3   port13   10000M   Yes
all    0-3   port14   10000M   Yes
all    0-3   port10   10000M   Yes
```

all	0-3	port9	10000M	Yes
all	0-3	port8	10000M	Yes
all	0-3	port7	10000M	Yes
all	0-3	port5	10000M	Yes
all	0-3	port6	10000M	Yes

FortiGate-3960E fast path architecture

The FortiGate-3960E features sixteen front panel 10GigE SFP+ interfaces and six 100GigE QSFP+ interfaces connected to sixteen NP6 processors through an Integrated Switch Fabric (ISF). Individual interfaces are not mapped to NP6 processors because of the integrated switch fabric. No special mapping is required for fastpath offloading or aggregate interfaces.



You can use the following `get` command to display the FortiGate-3960E NP6 configuration. The command output shows all NP6s connected to each interface (port) with cross-chip offloading supported for each port. You can also use the `diagnose npu np6 port-list` command to display this information.


```
get hardware npu np6 port-list
```

Chip	XAUI Ports	Max Speed	Cross-chip offloading
NP#0-7	0-3	port1	10000M Yes
NP#2	0-3	port2	10000M Yes
NP#0-7	0-3	port3	10000M Yes
NP#0-7	0-3	port4	10000M Yes
NP#0-7	0-3	port5	10000M Yes
NP#0-7	0-3	port6	10000M Yes
NP#0-7	0-3	port7	10000M Yes
NP#0-7	0-3	port8	10000M Yes
NP#0-7	0-3	port9	10000M Yes
NP#0-7	0-3	port10	10000M Yes
NP#0-7	0-3	port11	10000M Yes
NP#0-7	0-3	port12	10000M Yes
NP#0-7	0-3	port13	10000M Yes
NP#0-7	0-3	port14	10000M Yes
NP#0-7	0-3	port15	10000M Yes
NP#0-7	0-3	port16	10000M Yes
NP#7	0-3	port17	100000M Yes
NP#0-7	0-3	port18	100000M Yes
NP#10	0-3	port19	100000M Yes
NP#12-15	0-3	port20	100000M Yes
NP#8-15	0-3	port21	100000M Yes
NP#8-15	0-3	port22	100000M Yes

For information about optimizing FortiGate-3960E IPsec VPN performance, see [Optimizing FortiGate-3960E and 3980E IPsec VPN performance on page 1275](#).

FortiGate-3980E fast path architecture

The FortiGate-3980E features sixteen front panel 10GigE SFP+ interfaces and ten 100GigE QSFP28 interfaces connected to twenty-eight NP6 processors through an Integrated Switch Fabric (ISF). Individual interfaces are not mapped to NP6 processors because of the integrated switch fabric. No special mapping is required for fastpath offloading or aggregate interfaces.



You can use the following get command to display the FortiGate-3980E NP6 configuration. The command output shows all NP6s connected to each interface (port) with cross-chip offloading supported for each port. You can also use the `diagnose npu np6 port-list` command to display this information.

```
diagnose npu np6 port-list
Chip   XAUI Ports   Max   Cross-chip
        Speed   offloading
-----
NP#0-7      0-3 port1   10000M Yes
NP#0-7      0-3 port2   10000M Yes
NP#0-7      0-3 port3   10000M Yes
NP#0-7      0-3 port4   10000M Yes
NP#0-7      0-3 port5   10000M Yes
NP#0-7      0-3 port6   10000M Yes
NP#0-7      0-3 port7   10000M Yes
```

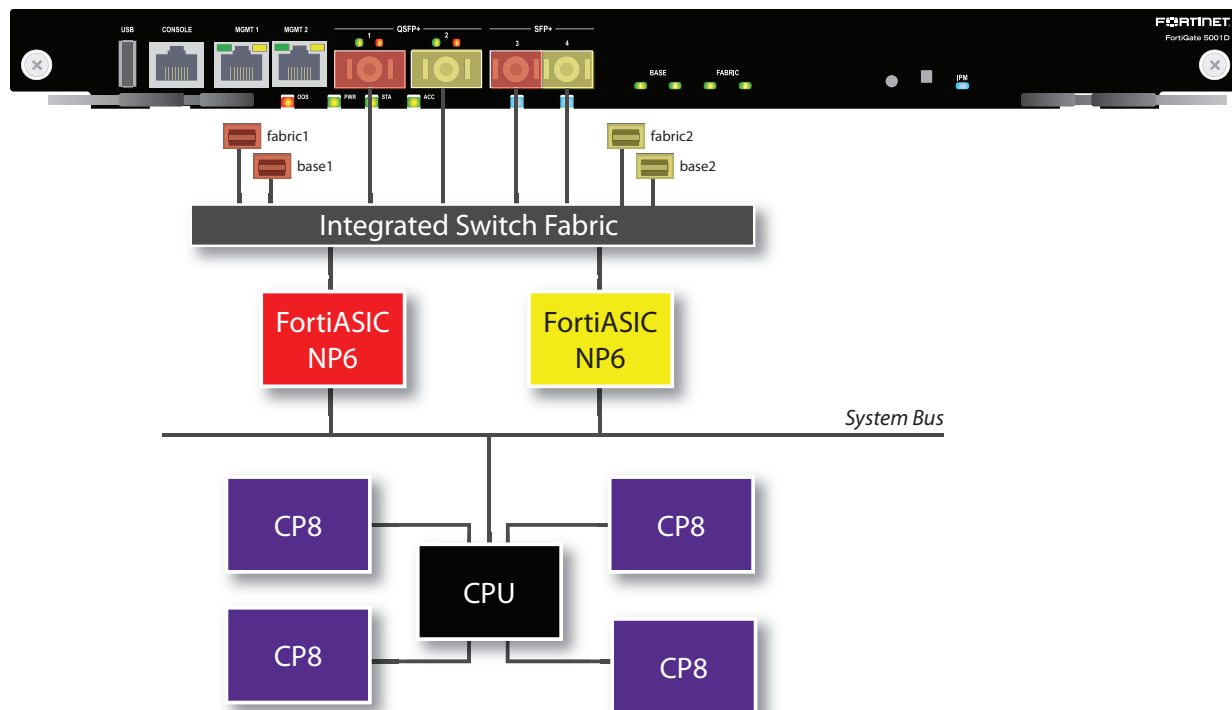
NP#0-7	0-3	port8	10000M	Yes
NP#0-7	0-3	port9	10000M	Yes
NP#0-7	0-3	port10	10000M	Yes
NP#0-7	0-3	port11	10000M	Yes
NP#0-7	0-3	port12	10000M	Yes
NP#0-7	0-3	port13	10000M	Yes
NP#0-7	0-3	port14	10000M	Yes
NP#0-7	0-3	port15	10000M	Yes
NP#0-7	0-3	port16	10000M	Yes
NP#0-7	0-3	port17	100000M	Yes
NP#0-7	0-3	port18	100000M	Yes
NP#8-27	0-3	port19	100000M	Yes
NP#8-27	0-3	port20	100000M	Yes
NP#8-27	0-3	port21	100000M	Yes
NP#8-27	0-3	port22	100000M	Yes
NP#8-27	0-3	port23	100000M	Yes
NP#8-27	0-3	port24	100000M	Yes
NP#8-27	0-3	port25	100000M	Yes
NP#8-27	0-3	port26	100000M	Yes

For information about optimizing FortiGate-3980E IPsec VPN performance, see [Optimizing FortiGate-3960E and 3980E IPsec VPN performance on page 1275](#).

FortiGate-5001D fast path architecture

The FortiGate5001D features two NP6 processors.

- port1, port3, fabric1 and base1 share connections to the first NP6 processor.
- port2, port4, fabric2 and base2 share connections to the second NP6 processor.



NP6 default interface mapping

You can use the following get command to display the FortiGate-5001D NP6 configuration. The command output shows two NP6s named NP6_0 and NP6_1. The output also shows the interfaces (ports) connected to each NP6. You can also use the `diagnose npu np6 port-list` command to display this information.

```
get hardware npu np6 port-list
```

Chip	XAUI	Ports	Max Speed	Cross-chip offloading
np6_0	0	port3	10G	Yes
	1			
	2	base1	1G	Yes
	3			
	0-3	port1	40G	Yes
	0-3	fabric1	40G	Yes
	0-3	fabric3	40G	Yes
	0-3	fabric5	40G	Yes
np6_1	0			
	1	port4	10G	Yes
	2			
	3	base2	1G	Yes
	0-3	port2	40G	Yes
	0-3	fabric2	40G	Yes
	0-3	fabric4	40G	Yes

NP6 interface mapping with split ports

If you use the following CLI command to split port1:

```
config system global
    set split-port port1
end
```

The new split ports (port1/1 to port 1/4) are mapped to the same NP6 as the port1 interface:

```
diagnose npu np6 port-list
```

Chip	XAUI	Ports	Max Speed	Cross-chip offloading
np6_0	0	port3	10G	Yes
	0	port1/1	10G	Yes
	1	port1/2	10G	Yes
	2	base1	1G	Yes
	2	port1/3	10G	Yes
	3	port1/4	10G	Yes
	0-3	fabric1	40G	Yes
	0-3	fabric3	40G	Yes
	0-3	fabric5	40G	Yes
np6_1	0			
	1	port4	10G	Yes
	2			

3	base2	1G	Yes
0-3	port2	40G	Yes
0-3	fabric2	40G	Yes
0-3	fabric4	40G	Yes

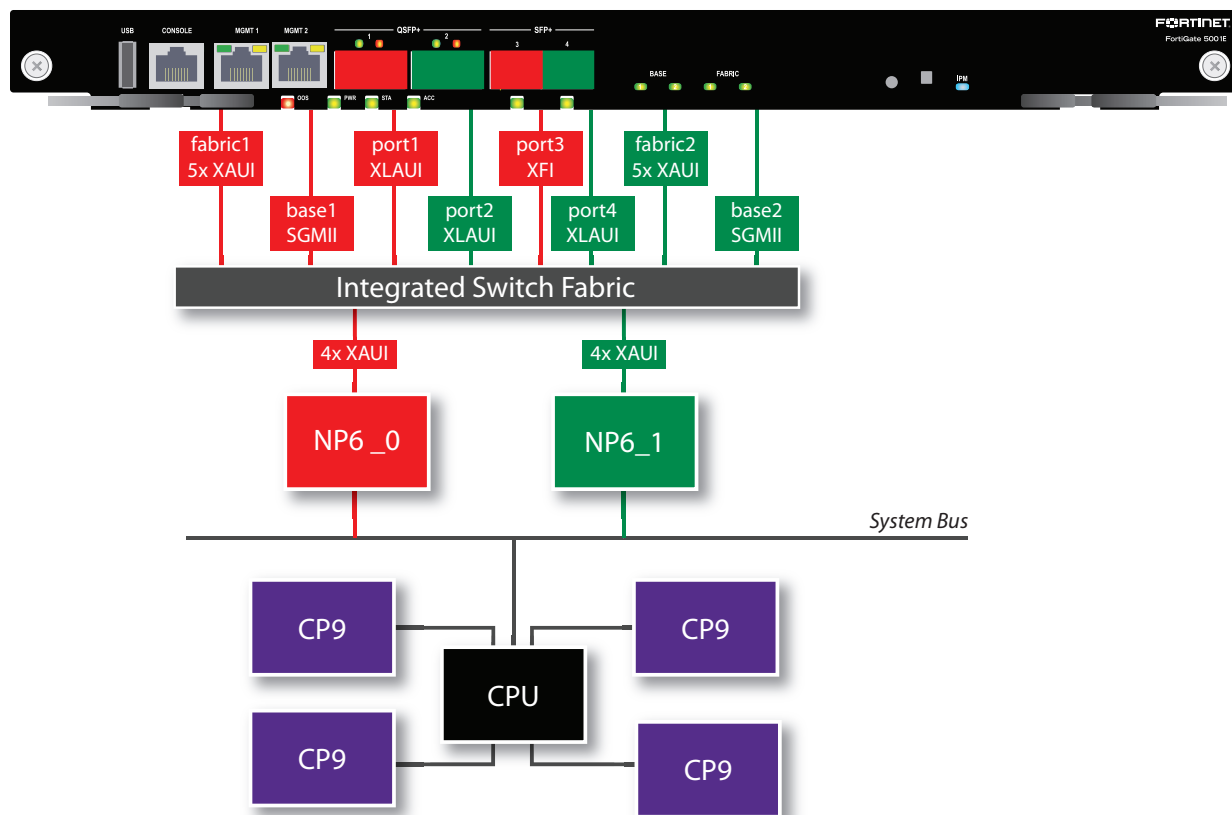
FortiGate-5001E and 5001E1 fast path architecture

The FortiGate-5001E and 5001E1 features two NP6 processors and an integrated switch fabric. The integrated switch fabric allows you to configure aggregate interfaces between interfaces connected to different NP6s and supports offloading between for traffic entering and exiting from any interfaces.

The NP6s are connected to network interfaces as follows:

- NP6_0 is connected to port1, port3, fabric1, and base1.
- NP6_1 is connected to port2, port4, fabric2, and base2.

The following diagram also shows the XAUI port connections between the NP6 processors and the front panel interfaces and the integrated switch fabric.



NP6 default interface mapping

You can use the following get command to display the FortiGate-5001E NP6 configuration. The command output shows two NP6s named NP6_0 and NP6_1. The output also shows the interfaces (ports) connected to each NP6.

You can also use the `diagnose npu np6 port-list` command to display this information.

```
get hardware npu np6 port-list
```

Chip	XAUI	Ports	Max Speed	Cross-chip offloading
np6_0	0	port3	10G	Yes
	1			
	2	base1	1G	Yes
	3			
	0-3	port1	40G	Yes
	0-3	fabric1	40G	Yes
	0-3	fabric3	40G	Yes
	0-3	fabric5	40G	Yes
np6_1	0			
	1	port4	10G	Yes
	2			
	3	base2	1G	Yes
	0-3	port2	40G	Yes
	0-3	fabric2	40G	Yes
	0-3	fabric4	40G	Yes

NP6 interface mapping with split ports

If you use the following CLI command to split port1:

```
config system global
  set split-port port1
end
```

The new split ports (port1/1 to port 1/4) are mapped to the same NP6 as the port1 interface:

```
diagnose npu np6 port-list
```

Chip	XAUI	Ports	Max Speed	Cross-chip offloading
np6_0	0	port3	10G	Yes
	0	port1/1	10G	Yes
	1	port1/2	10G	Yes
	2	base1	1G	Yes
	2	port1/3	10G	Yes
	3	port1/4	10G	Yes
	0-3	fabric1	40G	Yes
	0-3	fabric3	40G	Yes
	0-3	fabric5	40G	Yes
np6_1	0			
	1	port4	10G	Yes
	2			
	3	base2	1G	Yes
	0-3	port2	40G	Yes
	0-3	fabric2	40G	Yes

```

0-3 fabric4
-----
40G Yes
-----

```

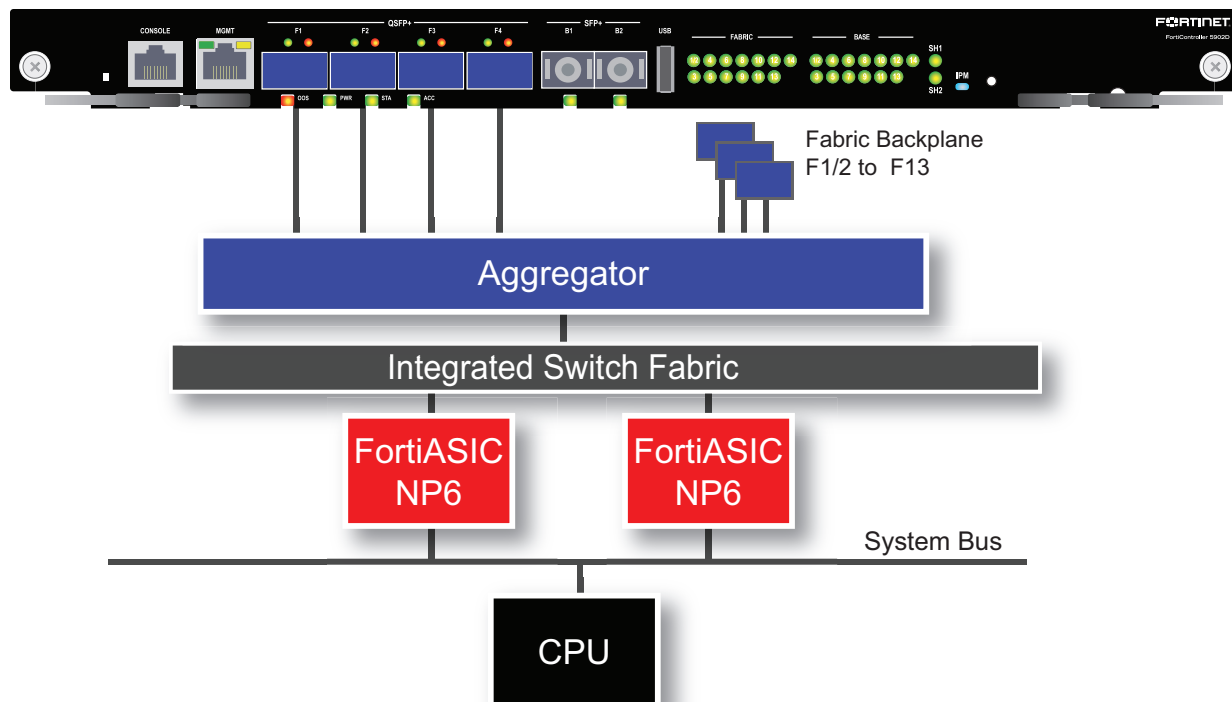
FortiController-5902D fast path architecture

The FortiController-5902D NP6 network processors and integrated switch fabric (ISF) provide hardware acceleration by offloading load balancing from the primary FortiController-5902D CPU. Network processors are especially useful for accelerating load balancing of TCP and UDP sessions.

The first packet of every new session is received by the primary FortiController-5902D and the primary FortiController-5902D uses its load balancing schedule to select the worker that will process the new session. This information is passed back to an NP6 network processor and all subsequent packets of the same sessions are offloaded to an NP6 network processor which sends the packet directly to a subordinate unit. Load balancing is effectively offloaded from the primary unit to the NP6 network processors resulting in a faster and more stable active-active cluster.

Traffic accepted by the FortiController-5902D F1 to F4 interfaces is that is processed by the primary FortiController-5902D is also be offloaded to the NP6 processors.

Individual FortiController-5902D interfaces are not mapped to NP6 processors. Instead an Aggregator connects the all fabric interfaces to the ISF and no special mapping is required for fastpath offloading.



NP6 content clustering mode interface mapping

FortiController-5902Ds run in content clustering mode and load balance sessions to FortiGate-5001D workers. Use the following command to enable content clustering:

```

config system elbc
    set mode content-cluster
    set inter-chassis-support enable
end

```

You can use the following get command to display the content clustering FortiController-5902D NP6 configuration. The output shows that all ports are mapped to all NP6 processors. You can also use the `diagnose npu np6 port-list` command to display this information.

```
get hardware npu np6 port-list
```

Chip	XAUI	Ports	Max Speed	Cross-chip offloading
all	0-3	f1	40000M	Yes
all	0-3	f2	40000M	Yes
all	0-3	f3	40000M	Yes
all	0-3	f4	40000M	Yes
all	0-3	np6_0_4	10000M	Yes
all	0-3	np6_0_5	10000M	Yes
all	0-3	elbc-ctrl/1-2	40000M	Yes
all	0-3	elbc-ctrl/3	40000M	Yes
all	0-3	elbc-ctrl/4	40000M	Yes
all	0-3	elbc-ctrl/5	40000M	Yes
all	0-3	elbc-ctrl/6	40000M	Yes
all	0-3	elbc-ctrl/7	40000M	Yes
all	0-3	elbc-ctrl/8	40000M	Yes
all	0-3	elbc-ctrl/9	40000M	Yes
all	0-3	elbc-ctrl/10	40000M	Yes
all	0-3	elbc-ctrl/11	40000M	Yes
all	0-3	elbc-ctrl/12	40000M	Yes
all	0-3	elbc-ctrl/13	40000M	Yes
all	0-3	elbc-ctrl/14	40000M	Yes

NP6 default interface mapping

You can use the following command to display the default FortiController-5902D NP6 configuration.

```
diagnose npu np6 port-list
```

Chip	XAUI	Ports	Max Speed	Cross-chip offloading
all	0-3	f1	40000M	Yes
all	0-3	f2	40000M	Yes
all	0-3	f3	40000M	Yes
all	0-3	f4	40000M	Yes
all	0-3	np6_0_4	10000M	Yes
all	0-3	np6_0_5	10000M	Yes
all	0-3	fabric1/2	40000M	Yes
all	0-3	fabric3	40000M	Yes
all	0-3	fabric4	40000M	Yes
all	0-3	fabric5	40000M	Yes
all	0-3	fabric6	40000M	Yes
all	0-3	fabric7	40000M	Yes

all	0-3	fabric8	40000M	Yes
all	0-3	fabric9	40000M	Yes
all	0-3	fabric10	40000M	Yes
all	0-3	fabric11	40000M	Yes
all	0-3	fabric12	40000M	Yes
all	0-3	fabric13	40000M	Yes
all	0-3	fabric14	40000M	Yes

FortiGate NP6lite architectures

This chapter shows the NP6lite architecture for the all FortiGate models that include NP6lite processors.

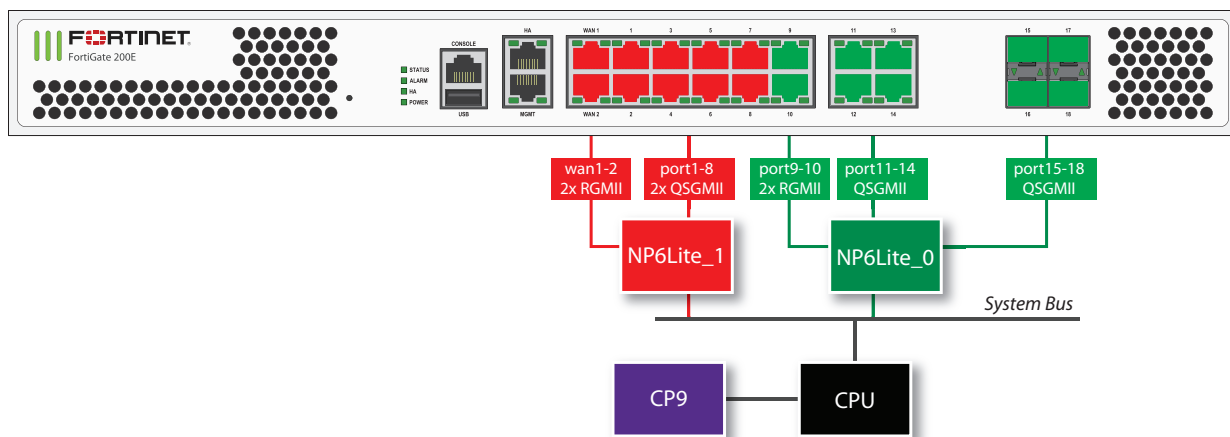
FortiGate-200E and 201E fast path architecture

The FortiGate-200E and 201E include two NP6lite processors. Because of this model does not include a switch fabric, you cannot create Link Aggregation Groups (LAGs) between interfaces connected to different NP6lites. As well traffic will only be offloaded if it enters and exits the FortiGate on interfaces connected to the same NP6lite.

The NP6lites are connected to network interfaces as follows:

- NP6lite_0 is connected to six 1GE RJ-45 interfaces (port9-port14) and four 1GE SFP interfaces (port15-18).
- NP6lite_1 is connected to ten 1GE RJ45 interfaces (wan1, wan2, port1-port8).

The following diagram also shows the RGMII and QSGMII port connections between the NP6lite processors and the front panel interfaces. Both RGMII and QSGMII interfaces operate at 1000Mbps. However, QSGMII interfaces can also negotiate to operate at lower speeds: 10, 100, and 1000Mbps. To connect the FortiGate-200E to networks with speeds lower than 1000Mbps use the QSGMII interfaces (port1-8 and port11-18).



You can use the following get command to display the FortiGate-200E or 201E NP6lite configuration. You can also use the `diagnose npu np6lite port-list` command to display this information.

```

get hardware npu np6lite port-list
Chip    XAUI Ports          Max    Cross-chip
        -----          Speed offloading
np6lite_0
  2    port9             1000M          NO
  1    port10            1000M          NO
  4    port11            1000M          NO
  3    port12            1000M          NO
  6    port13            1000M          NO
  5    port14            1000M          NO
  9    port15            1000M          NO
 10    port16            1000M          NO
  8    port17            1000M          NO
  7    port18            1000M          NO
np6lite_1
  2    wan1              1000M          NO
  1    wan2              1000M          NO
  4    port1             1000M          NO
  3    port2             1000M          NO
  6    port3             1000M          NO
  5    port4             1000M          NO
  8    port5             1000M          NO
  7    port6             1000M          NO
 10    port7             1000M          NO
  9    port8             1000M          NO

```

NP4 and NP4Lite acceleration

NP4 network processors provide fastpath acceleration by offloading communication sessions from the FortiGate CPU. When the first packet of a new session is received by an interface connected to an NP4 processor, just like any session connecting with any FortiGate interface, the session is forwarded to the FortiGate CPU where it is matched with a security policy. If the session is accepted by a security policy and if the session can be offloaded its session key is copied to the NP4 processor that received the packet. All of the rest of the packets in the session are intercepted by the NP4 processor and fast-pathed out of the FortiGate unit to their destination without ever passing through the FortiGate CPU. The result is enhanced network performance provided by the NP4 processor plus the network processing load is removed from the CPU. In addition, the NP4 processor can handle some CPU intensive tasks, like IPsec VPN encryption/decryption.



NP4lite processors have the same architecture and function in the same way as NP4 processors. All of the descriptions of NP4 processors in this document can be applied to NP4lite processors except where noted.

Session keys (and IPsec SA keys) are stored in the memory of the NP4 processor that is connected to the interface that received the packet that started the session. All sessions are fast-pathed and accelerated, even if they exit the FortiGate unit through an interface connected to another NP4. The key to making this possible is the Integrated Switch Fabric (ISF) that connects the NP4s and the FortiGate unit interfaces together. The ISF allows any port connectivity. All ports and NP4s can communicate with each other over the ISF.

There are no special ingress and egress fast path requirements because traffic enters and exits on interfaces connected to the same ISF. Most FortiGate models with multiple NP4 processors connect all interfaces and NP4 processors to the same ISF (except management interfaces) so this should not ever be a problem.

There is one limitation to keep in mind; the capacity of each NP4 processor. An individual NP4 processor has a capacity of 20 Gbps (10 Gbps ingress and 10 Gbps egress). Once an NP4 processor hits its limit, sessions that are over the limit are sent to the CPU. You can avoid this problem by as much as possible distributing incoming sessions evenly among the NP4 processors. To be able to do this you need to be aware of which interfaces connect to which NP4 processors and distribute incoming traffic accordingly.

Some FortiGate units contain one NP4 processor with all interfaces connected to it and to the ISF. As a result, offloading is supported for traffic between any pair of interfaces.

Some FortiGate units include NP4Lite processors. These network processors have the same functionality and limitations as NP4 processors but with about half the performance. NP4lite processors can be found in mid-range FortiGate models such as the FortiGate-200D and 240D.

Viewing your FortiGate NP4 processor configuration

To list the NP4 network processors on your FortiGate unit, use the following CLI command.

```
get hardware npu np4 list
```

The output lists the interfaces that have NP4 processors. For example, for a FortiGate-5001C:

```
get hardware npu np4 list
ID      Model      Slot      Interface
0       On-board
        port1 port2 port3 port4
        fabric1 base1 npu0-vlink0 npu0-vlink1
```

```
1      On-board                port5 port6 port7 port8
                                fabric2 base2 npu1-vlink0 npu1-vlink1
```

NP4lite CLI commands (disabling NP4Lite offloading)

If your FortiGate unit includes an NP4Lite processor the following commands will be available:

- Use the following command to disable or enable NP4Lite offloading. By default NP4lite offloading is enabled. If you want to disable NP4Lite offloading to diagnose a problem enter:

```
diagnose npu nplite fastpath disable
```

This command disables NP4Lite offloading until your FortiGate reboots. You can also re-enable offloading by entering the following command:

```
diagnose npu nplite fastpath enable
```

- NP4lite debug command. Use the following command to debug NP4Lite operation:

```
diagnose npl npl_debug {<parameters>}
```

NP4Lite option to disable offloading ICMP traffic in IPsec tunnels

In some cases ICMP traffic in IPsec VPN tunnels may be dropped by the NP4Lite processor due to a bug with the NP4Lite firmware. You can use the following command to avoid this problem by preventing the NP4Lite processor from offloading ICMP sessions in IPsec VPN tunnels. This command is only available on FortiGate models with NP4Lite processors, such as the FortiGate/FortiWiFi-60D.

```
config system npu
    set process-icmp-by-host {disable | enable}
end
```

The option is disabled by default and all ICMP traffic in IPsec VPN tunnels is offloaded where possible. If you are noticing that ICMP packets in IPsec VPN tunnels are being dropped you can disable this option and have all ICMP traffic processed by the CPU and not offloaded to the NP4Lite.

NP4 and NP4Lite processors and sFlow and NetFlow

Configuring sFlow or NetFlow on any interface disables all NP4 or NP4Lite offloading for all traffic on that interface.

Configuring NP4 traffic offloading

Offloading traffic to a network processor requires that the FortiGate unit configuration and the traffic itself is suited to hardware acceleration. There are requirements for path the sessions and the individual packets.

NP4 session fast path requirements

Sessions must be fast path ready. Fast path ready session characteristics are:

- Layer 2 type/length must be 0x0800 (IEEE 802.1q VLAN specification is supported)
- Layer 3 protocol must be IPv4
- Layer 4 protocol must be UDP, TCP or ICMP
- Layer 3 / Layer 4 header or content modification must not require a session helper (for example, SNAT, DNAT, and TTL reduction are supported, but application layer content modification is not supported)

- Firewall policies must not include proxy-based security features (proxy-based virus scanning, proxy-based web filtering, DNS filtering, DLP, Anti-Spam, VoIP, ICAP, Web Application Firewall, or Proxy options).
- If the FortiGate supports NTurbo, firewall policies can include flow-based security features (IPS, Application Control CASI, flow-based antivirus, or flow-based web filtering) .
- Origin must not be local host (the FortiGate unit)



If you disable anomaly checks by Intrusion Prevention (IPS), you can still enable NP4 hardware accelerated anomaly checks using the `fp-anomaly` field of the `config system interface` CLI command. See [Offloading NP4 anomaly detection on page 1336](#)

If a session is not fast path ready, the FortiGate unit will not send the session key to the network processor(s). Without the session key, all session key lookup by a network processor for incoming packets of that session fails, causing all session packets to be sent to the FortiGate unit's main processing resources, and processed at normal speeds.

If a session is fast path ready, the FortiGate unit will send the session key to the network processor(s). Session key lookup then succeeds for subsequent packets from the known session.

Packet fast path requirements

Packets within the session must then also meet packet requirements.

- Incoming packets must not be fragmented.
- Outgoing packets must not require fragmentation to a size less than 385 bytes. Because of this requirement, the configured MTU (Maximum Transmission Unit) for network processors' network interfaces must also meet or exceed the network processors' supported minimum MTU of 385 bytes.

If packet requirements are not met, an individual packet will use FortiGate unit main processing resources, regardless of whether other packets in the session are offloaded to the specialized network processor(s).

In some cases, due to these requirements, a protocol's session(s) may receive a mixture of offloaded and non-offloaded processing.

For example, FTP uses two connections: a control connection and a data connection. The control connection requires a session helper, and cannot be offloaded, but the data connection does not require a session helper, and can be offloaded. Within the offloadable data session, fragmented packets will not be offloaded, but other packets will be offloaded.

Some traffic types differ from general offloading requirements, but still utilize some of the network processors' encryption and other capabilities. Exceptions include IPsec traffic and active-active high availability (HA) load balanced traffic.

Mixing fast path and non-fast path traffic

If packet requirements are not met, an individual packet will be processed by the FortiGate CPU regardless of whether other packets in the session are offloaded to the NP4.

Also, in some cases, a protocol's session(s) may receive a mixture of offloaded and non-offloaded processing. For example, VoIP control packets may not be offloaded but VoIP data packets (voice packets) may be offloaded.

Disabling NP4 and NP4lite hardware acceleration (fastpath)

You can use the following command to disable NP4 offloading for all traffic. This option disables NP4 offloading for all traffic for all NP4 and NP4lite processors.

```
config system npu
    set fastpath disable
end
```

Increasing NP4 offloading capacity using link aggregation groups (LAGs)

NP4 processors can offload sessions received by interfaces in link aggregation groups (LAGs) (IEEE 802.3ad). A LAG combines more than one physical interface into a group that functions like a single interface with a higher capacity than a single physical interface. For example, you could use a LAG if you want to offload sessions on a 3 Gbps link by adding three 1Gbps interfaces to the same LAG.

All offloaded traffic types are supported by LAGs, including IPsec VPN traffic. Just like with normal interfaces, traffic accepted by a LAG is offloaded by the NP4 processor connected to the interfaces in the LAG that receive the traffic to be offloaded. If all interfaces in a LAG are connected to the same NP4 processor, traffic received by that LAG is offloaded by that NP4 processor. The amount of traffic that can be offloaded is limited by the capacity of the NP4 processor.

If a FortiGate has two or more NP4 processors connected by an integrated switch fabric (ISF), you can use LAGs to increase offloading by sharing the traffic load across multiple NP4 processors. You do this by adding physical interfaces connected to different NP4 processors to the same LAG.

Adding a second NP4 processor to a LAG effectively doubles the offloading capacity of the LAG. Adding a third further increases offloading. The actual increase in offloading capacity may not actually be doubled by adding a second NP4 or tripled by adding a third. Traffic and load conditions and other factors may limit the actual offloading result.

The increase in offloading capacity offered by LAGs and multiple NP4s is supported by the ISF that allows multiple NP4 processors to share session information. On models that have more than one NP4 and no ISF, if you attempt to add interfaces connected to different NP4 processors to a LAG the system displays an error message.

There are also a few limitations to LAG NP4 offloading support for IPsec VPN:

- IPsec VPN anti-replay protection cannot be used if IPsec is configured on a LAG that has interfaces connected to multiple NP4 processors.
- Using a LAG connected to multiple NP4 processors for decrypting incoming IPsec VPN traffic may cause some of the incoming traffic to be decrypted by the CPU. So this configuration is not recommended since not all decryption is offloaded. (Using a LAG connected to multiple NP4 processors for encrypting outgoing IPsec VPN traffic is supported with no limitations.)
- Because the encrypted traffic for one IPsec VPN tunnel has the same 5-tuple, the traffic from one tunnel can only can be balanced to one interface in a LAG. This limits the maximum throughput for one IPsec VPN tunnel in an NP4 LAG group to 1Gbps.

NP4 traffic shaping offloading

Accelerated Traffic shaping is supported with the following limitations.

- NP4 processors support policy-based traffic shaping. However, fast path traffic and traffic handled by the FortiGate CPU (slow path) are controlled separately, which means the policy setting on fast path does not consider the traffic on the slow path.
- The port based traffic policing as defined by the inbandwidth and outbandwidth CLI commands is not supported.
- DSCP configurations are supported.
- Per-IP traffic shaping is supported.
- QoS in general is not supported.

You can also use the traffic shaping features of the FortiGate unit's main processing resources by disabling NP4 offloading. See [Disabling NP offloading for firewall policies on page 1250](#).

NP4 IPsec VPN offloading

NP4 processors improve IPsec tunnel performance by offloading IPsec encryption and decryption.

Requirements for hardware accelerated IPsec encryption or decryption are a modification of general offloading requirements. Differing characteristics are:

- Origin can be local host (the FortiGate unit)
- In Phase 1 configuration, Local Gateway IP must be specified as an IP address of a network interface for a port attached to a network processor
- SA must have been received by the network processor
- in Phase 2 configuration:
 - encryption algorithm must be DES, 3DES, AES-128, AES-192, AES-256, or null
 - authentication must be MD5, SHA1, or null
 - if encryption is null, authentication must not also be null

To apply hardware accelerated encryption and decryption, the FortiGate unit's main processing resources must first perform Phase 1 negotiations to establish the security association (SA). The SA includes cryptographic processing instructions required by the network processor, such as which encryption algorithms must be applied to the tunnel. After ISAKMP negotiations, the FortiGate unit's main processing resources send the SA to the network processor, enabling the network processor to apply the negotiated hardware accelerated encryption or decryption to tunnel traffic.

Possible accelerated cryptographic paths are:

- IPsec decryption offload
 - Ingress ESP packet > Offloaded decryption > Decrypted packet egress (fast path)
 - Ingress ESP packet > Offloaded decryption > Decrypted packet to FortiGate unit's main processing resources
- IPsec encryption offload
 - Ingress packet > Offloaded encryption > Encrypted (ESP) packet egress (fast path)
 - Packet from FortiGate unit's main processing resources > Offloaded encryption > Encrypted (ESP) packet egress

Configuring inter-VDOM link acceleration with NP4 processors

FortiGate units with NP4 processors include inter-VDOM links that can be used to accelerate inter-VDOM link traffic.



Traffic is blocked if you enable IPS for traffic passing over inter-VDOM links if that traffic is being offloaded by an NP4 processor. If you disable NP4 offloading traffic will be allowed to flow. You can disable offloading in individual firewall policies by disabling `auto-asic-offload` for those policies. You can also use the following command to disable all IPS offloading

```
config ips global
    set np-accel-mode none
    set cp-accel-mode none
end
```

- For a FortiGate unit with two NP4 processors there are also two inter-VDOM links, each with two interfaces:

- **npu0-vlink:**

npu0-vlink0
npu0-vlink1

- **npu1-vlink:**

npu1-vlink0
npu1-vlink1

These interfaces are visible from the GUI and CLI. For a FortiGate unit with NP4 interfaces, enter the following CLI command (output shown for a FortiGate-5001B):

```
get hardware npu np4 list
ID      Model      Slot      Interface
0       On-board
        port1 port2 port3 port4
        fabric1 base1 npu0-vlink0 npu0-vlink1
1       On-board
        port5 port6 port7 port8
        fabric2 base2 npu1-vlink0 npu1-vlink1
```

By default the interfaces in each inter-VDOM link are assigned to the root VDOM. To use these interfaces to accelerate inter-VDOM link traffic, assign each interface in a pair to the VDOMs that you want to offload traffic between. For example, if you have added a VDOM named New-VDOM to a FortiGate unit with NP4 processors, you can go to **System > Network > Interfaces** and edit the **npu0-vlink1** interface and set the **Virtual Domain** to **New-VDOM**.

This results in an inter-VDOM link between root and New-VDOM. You can also do this from the CLI:

```
config system interface
  edit npu0-vlink1
    set vdom New-VDOM
  end
```

Using VLANs to add more accelerated inter-VDOM links

You can add VLAN interfaces to the accelerated inter-VDOM links to create inter-VDOM links between more VDOMs. For the links to work, the VLAN interfaces must be added to the same inter-VDOM link, must be on the same subnet, and must have the same VLAN ID.

For example, to accelerate inter-VDOM link traffic between VDOMs named Marketing and Engineering using VLANs with VLAN ID 100 go to **System > Network > Interfaces** and select **Create New** to create the VLAN interface associated with the Marketing VDOM:

Name	Marketing-link
Type	VLAN
Interface	npu0-vlink0
VLAN ID	100
Virtual Domain	Marketing
IP/Network Mask	172.20.120.12/24

Create the inter-VDOM link associated with Engineering VDOM:

Name	Engineering-link
Type	VLAN
Interface	npu0-vlink1
VLAN ID	100
Virtual Domain	Engineering
IP/Network Mask	172.20.120.22/24

Or do the same from the CLI:

```

config system interface
  edit Marketing-link
    set vdom Marketing
    set ip 172.20.120.12/24
    set interface npu0-vlink0
    set vlanid 100
  next
  edit Engineering-link
    set vdom Engineering
    set ip 172.20.120.22/24
    set interface npu0-vlink1
    set vlanid 100

```

Confirm that the traffic is accelerated

Use the following CLI commands to obtain the interface index and then correlate them with the session entries. In the following example traffic was flowing between new accelerated inter-VDOM links and physical ports port1 and port 2 also attached to the NP4 processor.

diagnose ip address list

```

IP=172.31.17.76->172.31.17.76/255.255.252.0 index=5 devname=port1
IP=10.74.1.76->10.74.1.76/255.255.252.0 index=6 devname=port2
IP=172.20.120.12->172.20.120.12/255.255.255.0 index=55 devname=IVL-VLAN1_ROOT
IP=172.20.120.22->172.20.120.22/255.255.255.0 index=56 devname=IVL-VLAN1_VDOM1

```

diagnose sys session list

```

session info: proto=1 proto_state=00 duration=282 expire=24 timeout=0 session info:
  proto=1 proto_state=00 duration=124 expire=59 timeout=0 flags=00000000
  sockflag=00000000 sockport=0 av_idx=0 use=3
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 policy_dir=0 tunnel=/
state=may_dirty npu
statistic(bytes/packets/allow_err): org=180/3/1 reply=120/2/1 tuples=2
origin->sink: org pre->post, reply pre->post dev=55->5/5->55
  gwy=172.31.19.254/172.20.120.22
hook=post dir=org act=snat 10.74.2.87:768->10.2.2.2:8(172.31.17.76:62464)
hook=pre dir=reply act=dnat 10.2.2.2:62464->172.31.17.76:0(10.74.2.87:768)
misc=0 policy_id=4 id_policy_id=0 auth_info=0 chk_client_info=0 vd=0
serial=0000004e tos=ff/ff ips_view=0 app_list=0 app=0
dd_type=0 dd_mode=0
per_ip_bandwidth meter: addr=10.74.2.87, bps=880
npu_state=00000000
npu info: flag=0x81/0x81, offload=4/4, ips_offload=0/0, epid=160/218, ipid=218/160,
vlan=32769/0

session info: proto=1 proto_state=00 duration=124 expire=20 timeout=0 flags=00000000
  sockflag=00000000 sockport=0 av_idx=0 use=3
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 policy_dir=0 tunnel=/
state=may_dirty npu
statistic(bytes/packets/allow_err): org=180/3/1 reply=120/2/1 tuples=2
origin->sink: org pre->post, reply pre->post dev=6->56/56->6
  gwy=172.20.120.12/10.74.2.87
hook=pre dir=org act=noop 10.74.2.87:768->10.2.2.2:8(0.0.0.0:0)

```

```
hook=post dir=reply act=noop 10.2.2.2:768->10.74.2.87:0(0.0.0.0:0)
misc=0 policy_id=3 id_policy_id=0 auth_info=0 chk_client_info=0 vd=1
serial=0000004d tos=ff/ff ips_view=0 app_list=0 app=0
dd_type=0 dd_mode=0
per_ip_bandwidth meter: addr=10.74.2.87, bps=880
npu_state=00000000
npu info: flag=0x81/0x81, offload=4/4, ips_offload=0/0, epid=219/161, ipid=161/219,
vlan=0/32769
total session 2
```

Offloading NP4 anomaly detection

Network interfaces associated with a port attached to an NP4 processor can be configured to offload anomaly checking to the NP4 processor. This anomaly checking happens before other offloading and separately from DoS policy anomaly checking. Using the following command, each FortiGate interface can have a different anomaly checking configuration even if they are connected to the same NP4 processor.



The options available for this command apply anomaly checking for NP4 sessions in the same way as the command described in [Configuring individual NP6 processors on page 1265](#) applies anomaly checking for NP6 sessions.

```
config system interface
  edit <port-name>
    set fp-anomaly <anomalies>
  end
```

where <anomalies> can be one, more than one or all of the following:

Anomaly	Description
drop_icmp_frag	Drop ICMP fragments to pass.
drop_icmpland	Drop ICMP Land.
drop_ipland	Drop IP Land.
drop_iplsrr	Drop IP with Loose Source Record Route option.
drop_iprr	Drop IP with Record Route option.
drop_ipsecurity	Drop IP with Security option.
drop_ipssrr	Drop IP with Strict Source Record Route option.
drop_ipstream	Drop IP with Stream option.
drop_iptimestamp	Drop IP with Timestamp option.

Anomaly	Description
drop_ipunknown_option	Drop IP with malformed option.
drop_ipunknown_prot	Drop IP with Unknown protocol.
drop_tcp_fin_noack	Drop TCP FIN with no ACK flag set to pass.
drop_tcp_no_flag	Drop TCP with no flag set to pass.
drop_tcpland	Drop TCP Land.
drop_udpland	Drop UDP Land.
drop_winnuke	Drop TCP WinNuke.
pass_icmp_frag	Allow ICMP fragments to pass.
pass_icmpland	Allow ICMP Land to pass.
pass_ipland	Allow IP land to pass.
pass_iplsrr	Allow IP with Loose Source Record Route option to pass.
pass_iprr	Allow IP with Record Route option to pass.
pass_ipsecurity	Allow IP with Security option to pass.
pass_ipssrr	Allow IP with Strict Source Record Route option to pass.
pass_ipstream	Allow IP with Stream option to pass.
pass_iptimestamp	Allow IP with Timestamp option to pass.
pass_ipunknown_option	Allow IP with malformed option to pass.
pass_ipunknown_prot	Allow IP with Unknown protocol to pass.

Anomaly	Description
pass_tcp_fin_noack	Allow TCP FIN with no ACT flag set to pass.
pass_tcp_no_flag	Allow TCP with no flag set to pass.
pass_tcpland	Allow TCP Land to pass.
pass_udpland	Allow UDP Land to pass.
pass_winnuke	Allow TCP WinNuke to pass.

Example

You might configure an NP4 to drop packets with TCP WinNuke or unknown IP protocol anomalies, but to pass packets with an IP time stamp, using hardware acceleration provided by the network processor.

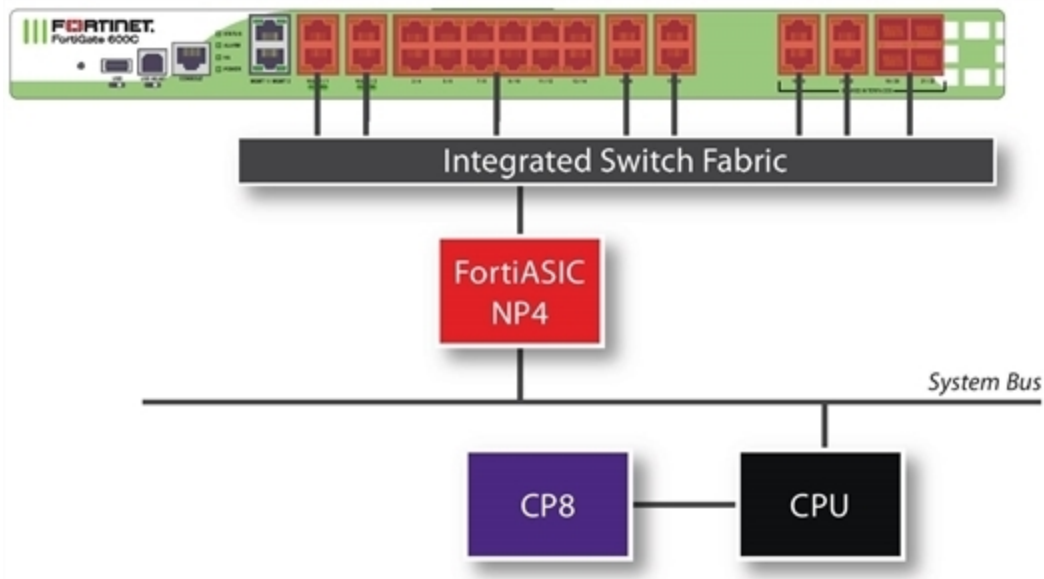
```
config system interface
  edit port1
    set fp-anomaly drop_winnuke drop_ipunknown_prot pass_iptimestamp
  end
```

FortiGate NP4 architectures

This chapter shows the NP4 architecture for the all FortiGate models that include NP4 processors.

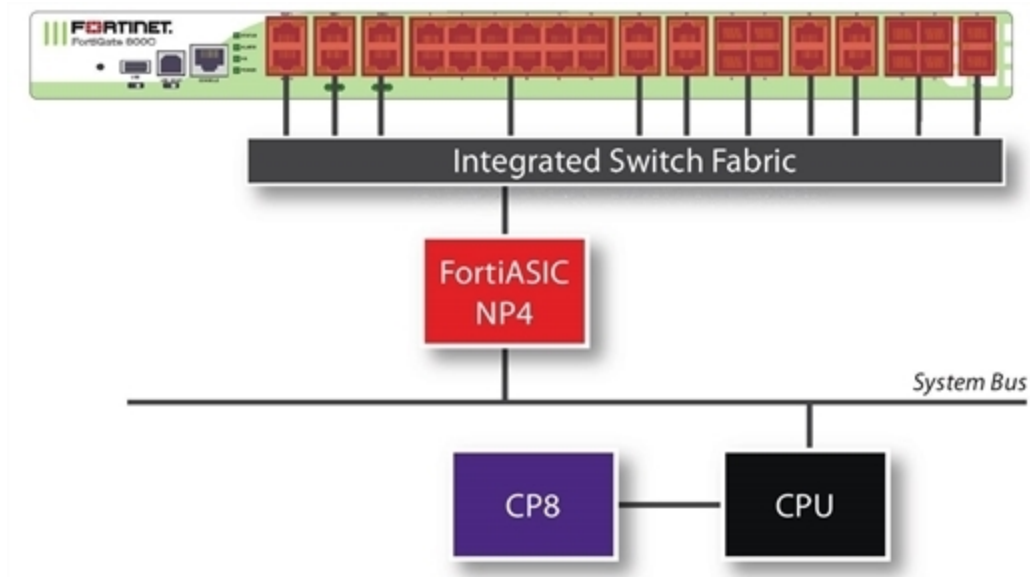
FortiGate-600C

The FortiGate-600C features one NP4 processor. All the ports are connected to this NP4 over the Integrated Switch Fabric. Port1 and port2 are dual failopen redundant RJ-45 ports. Port3-port22 are RJ-45 ethernet ports, and there are four 1Gb SFP interface ports duplicating the port19-port22 connections.




FortiGate-800C

The FortiGate-800C features one NP4 processor. All the ports are connected to this NP4. Port1 and port2 are dual failopen redundant RJ-45 ports. Port3-port22 are RJ-45 Ethernet ports, and there are eight 1Gb SFP interface ports duplicating the port15-18 and port19-port22 connections. There are also two 10Gb SFP+ ports, port23 and port24.



FortiGate-1000C

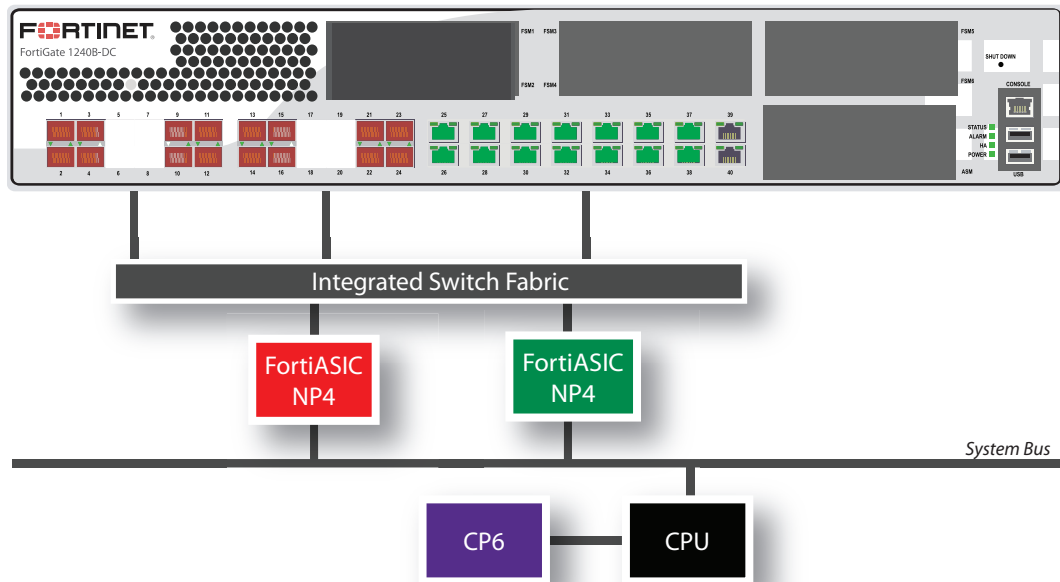
The FortiGate-1000C features one NP4 processor. All the ports are connected to this NP4. Port1 and port2 are dual failopen redundant RJ-45 ports. Port3-port22 are RJ-45 ethernet ports, and there are eight 1Gb SFP interface ports duplicating the port15-18 and port19-port22 connections. There are also two 10Gb SFP+ ports, port23 and port24.



FortiGate-1240B

The FortiGate-1240B features two NP4 processors:

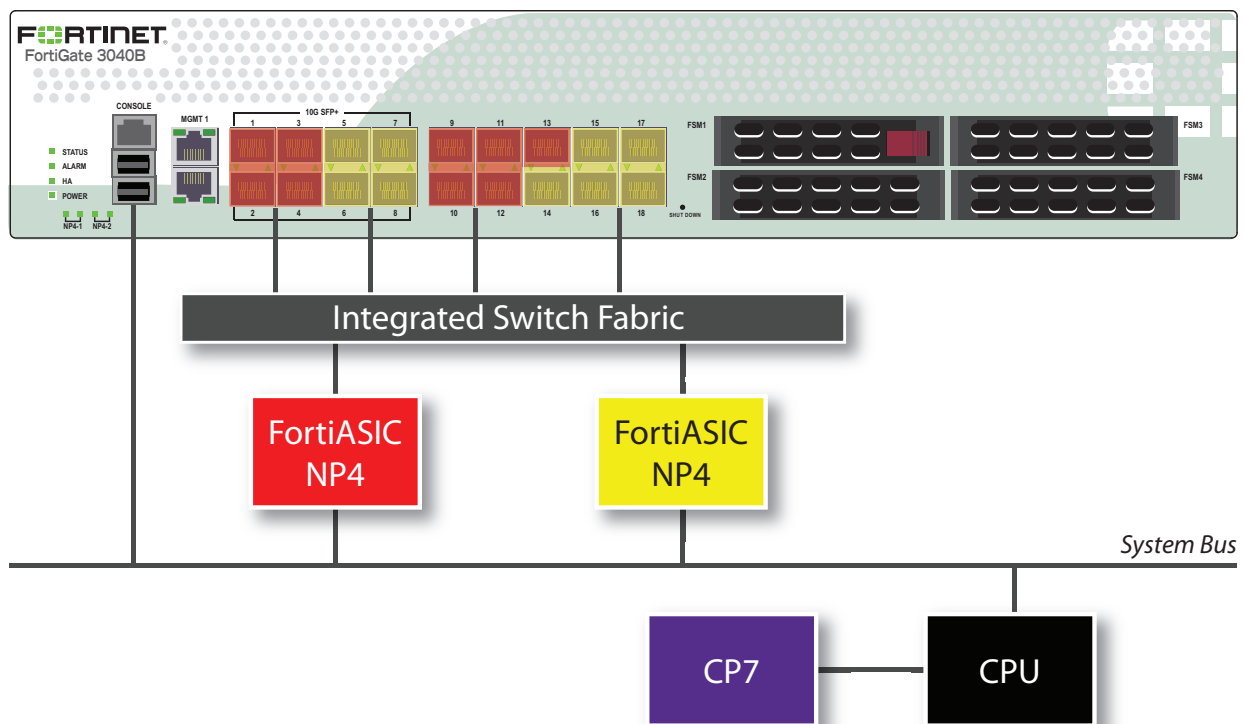
- Port1 to port24 are 1Gb SFP interfaces connected to one NP4 processor.
- Port25 to port38 are RJ-45 ethernet ports, connected to the other NP4 processor.
- Port39 and port40 are not connected to an NP4 processor.



FortiGate-3040B

The FortiGate-3040B features two NP4 processors:

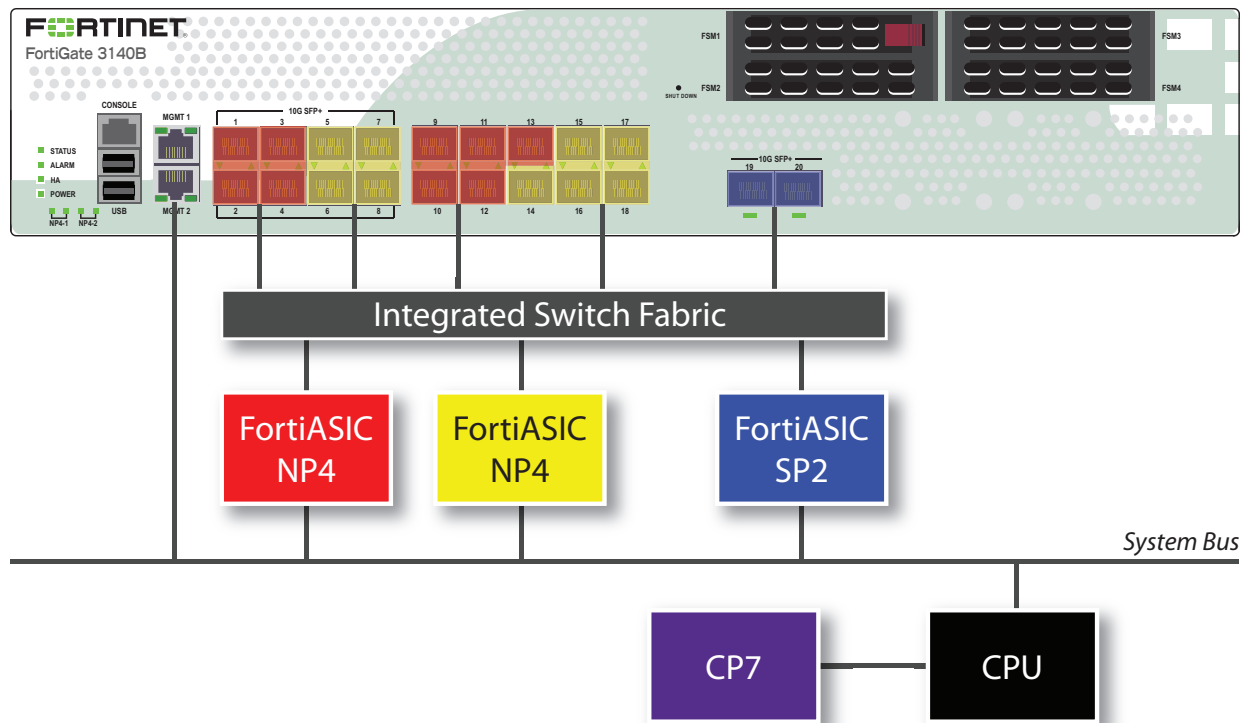
- The 10Gb interfaces, port1, port2, port3, port4, and the 1Gb interfaces, port9, port10, port11, port12, port13, share connections to one NP4 processor.
- The 10Gb interfaces, port5, port6, port7, port8, and the 1Gb interfaces, port14, port15, port16, port17, port18, share connections to the other NP4 processor.



FortiGate-3140B

The FortiGate-3140B features two NP4 processors and one SP2 processor:

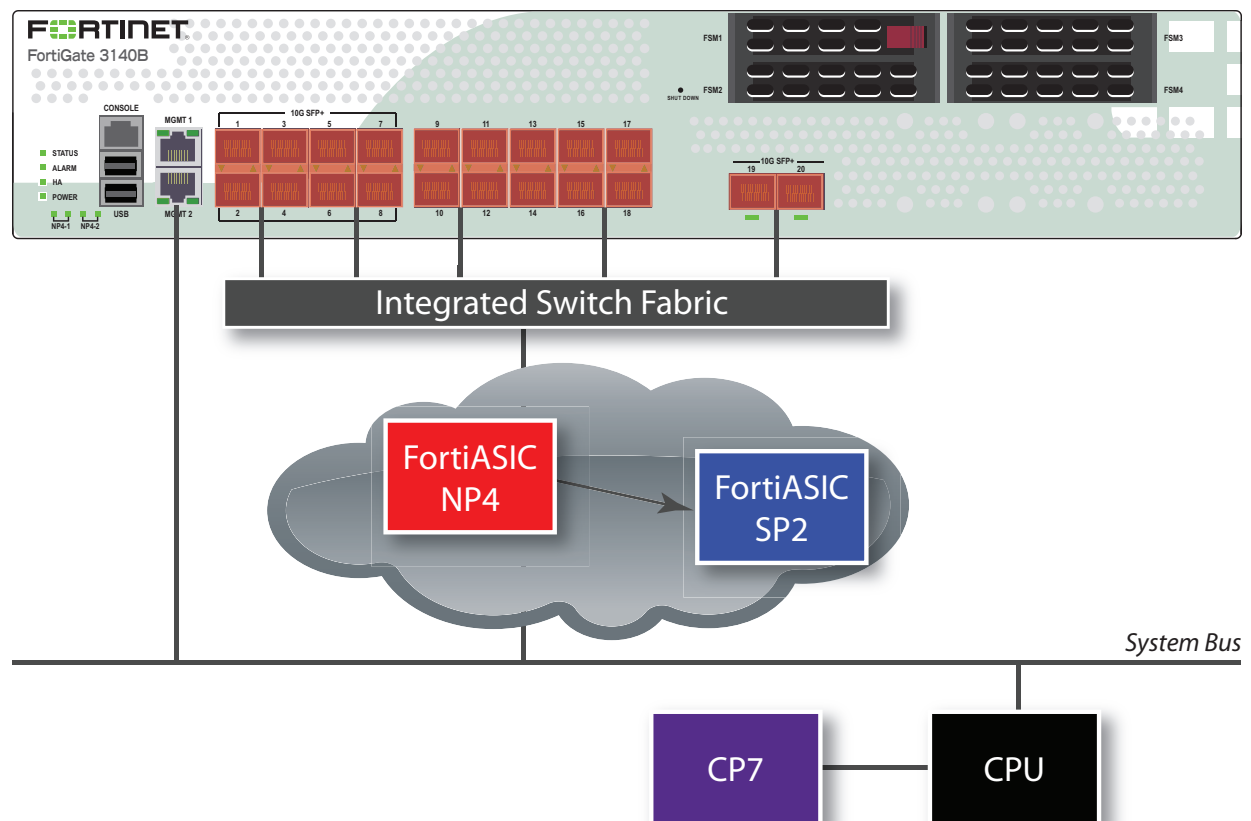
- The 10Gb interfaces, port1, port2, port3, port4, and the 1Gb interfaces, port9, port10, port11, port12, port13, share connections to one NP4 processor.
- The 10Gb interfaces, port5, port6, port7, port8, and the 1Gb interfaces, port14, port15, port16, port17, port18, share connections to the other NP4 processor.
- The 10Gb interfaces, port19 and port20, share connections to the SP2 processor.



FortiGate-3140B — load balance mode

The FortiGate-3140B load balance mode allows you increased flexibility in how you use the interfaces on the FortiGate unit. When enabled, traffic between any two interfaces (excluding management and console) is accelerated. Traffic is not limited to entering and leaving the FortiGate unit in specific interface groupings to benefit from NP4 and SP2 acceleration. You can use any pair of interfaces.

Security acceleration in this mode is limited, however. Only IPS scanning is accelerated in load balance mode.



To enable this feature, issue this CLI command.

```
config system global
    set sp-load-balance enable
end
```

The FortiGate unit will then restart.

To return to the default mode, issue this CLI command.

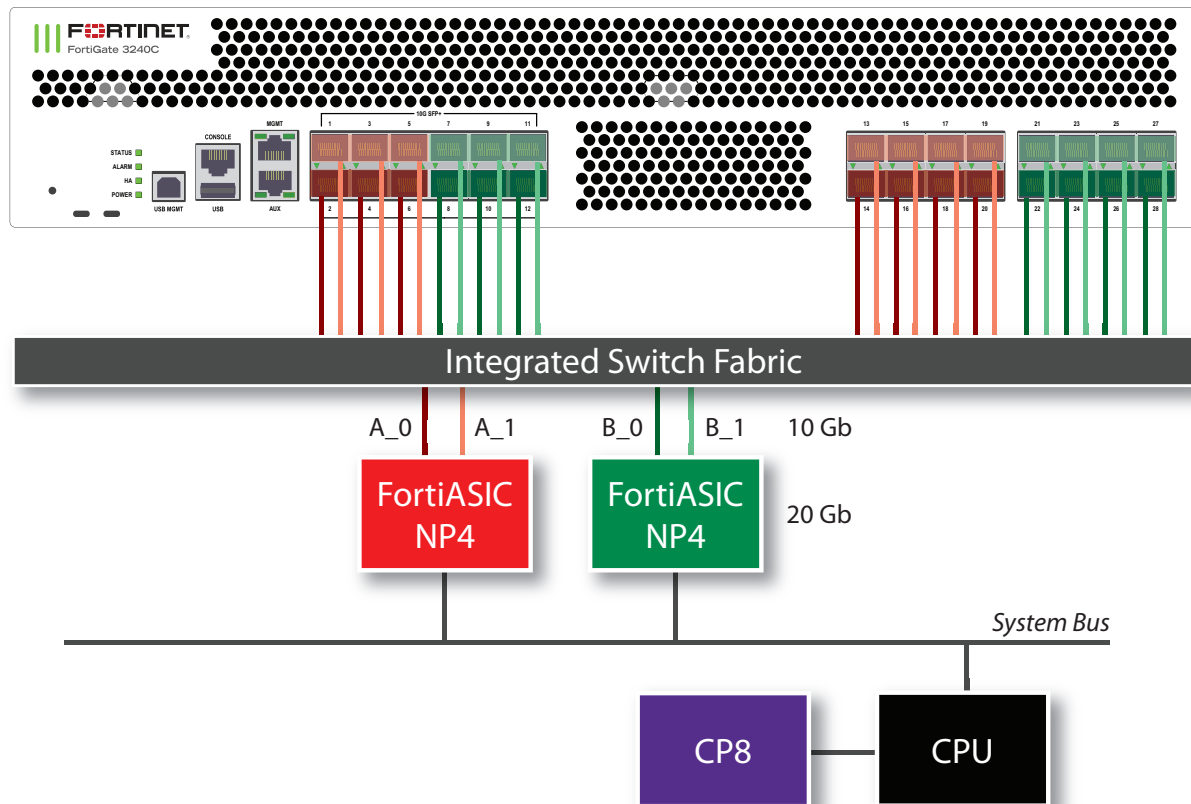
```
config system global
    set sp-load-balance disable
end
```

FortiGate-3240C

The FortiGate-3240C features two NP4 processors:

- The 10Gb interfaces, port1 through port6, and the 1Gb interfaces, port13 through port20, share connections to one NP4 processor.
- The 10Gb interfaces, port7 through port12, and the 1Gb interfaces, port21 through port28, share connections to the other NP4 processor.

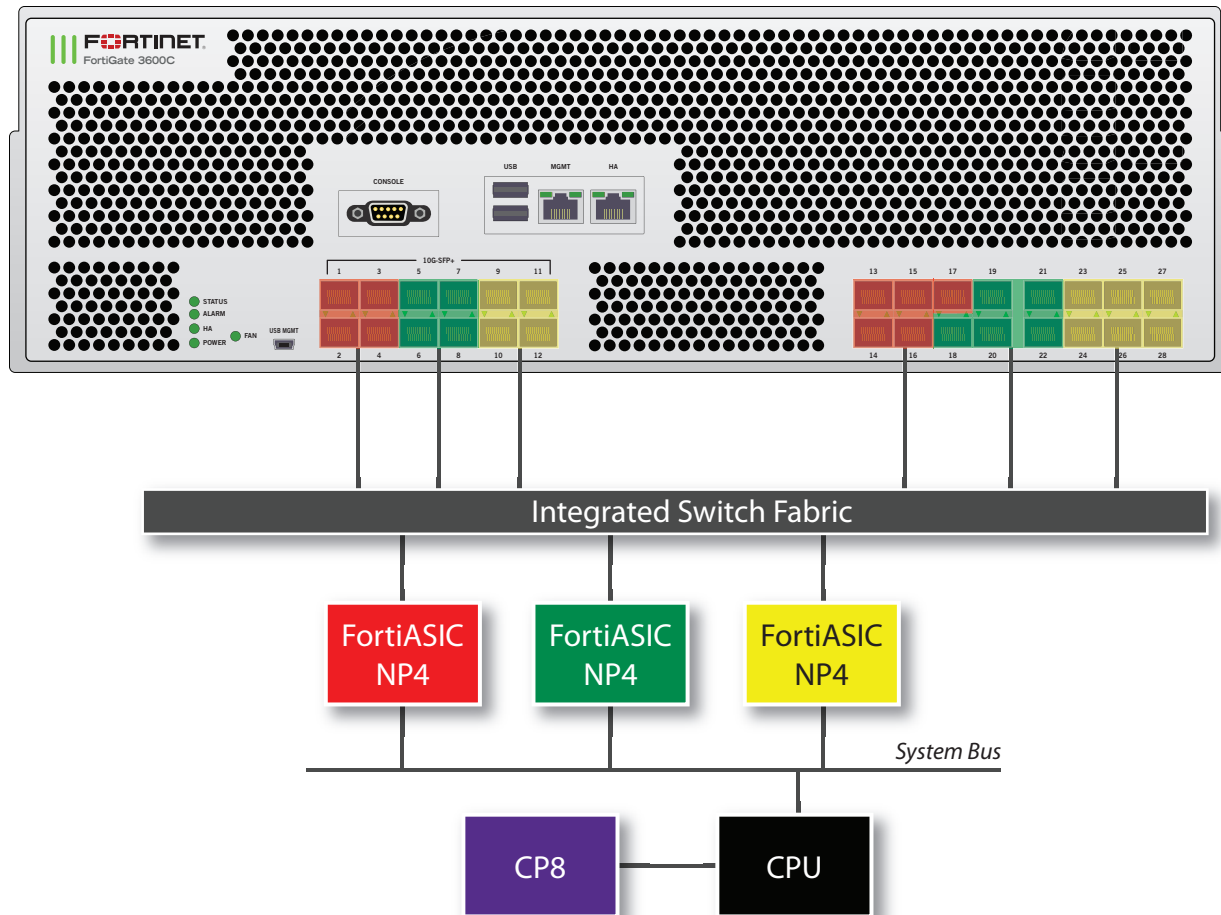
In addition to the ports being divided between the two NP4 processors, they are further divided between the two connections to each processor. Each NP4 can process 20 Gb of network traffic per second and each of two connections to each NP4 can move 10Gb of data to the processor per second, so the ideal configuration would have no more than 10 Gb of network traffic to each connection of each NP4 at any time.



FortiGate-3600C

The FortiGate-3600C features three NP4 processors:

- The 10Gb interfaces, port1-port4, and the 1Gb interfaces, port13-port17, share connections to one NP4 processor.
- The 10Gb interfaces, port5-port8, and the 1Gb interfaces, port18-port22 share connections to the second NP4 processor.
- The 10Gb interfaces, port9-port12, and the 1Gb interfaces, port23-port28 share connections to the third NP4 processor.



FortiGate-3600C XAUI links

The FortiGate-3600C uses XAUI links for communication between physical Ethernet ports and the integrated switch fabric.

Each XAUI link has a maximum bandwidth of 10-Gigabits. The reason you may need to know about the XAUI link in NP4 configurations is because of this 10-Gigabit limit. Because of this limitation, the total amount of data processed by all Ethernet interfaces connected to an XAUI link cannot exceed 10 gigabits. In some cases this may limit the amount of bandwidth that the FortiGate can process.

Each NP4 processor connects to the integrated switch fabric through two XAUI links: XAUI0 and XAUI1. All of the odd numbered Ethernet interfaces use XAUI0 and all of the even numbered interfaces use XAUI1:

NPU1

XAUI0 = port1, port3, port13, port15, port17

XAUI1 = port2, port4, port14, port16

NPU2

XAUI0 = port5, port7, port18, port20, port22

XAUI1 = port6, port8, port19, port21

NPU3

XAUI0 = port9, port11, port23, port25, port27

XAUI1 = port10, port12, port24, port26, port28

Usually you do not have to be concerned about XAUI link mapping. However, if a FortiGate-3600C NP4 interface is processing a very high amount of traffic you should distribute that traffic among both of the XAUI links connected to it. So if you have a very high volume of traffic flowing between two networks you should connect both networks to the same NP4 processor but to different XAUI links. For example, you could connect one network to Ethernet port5 and the other network to Ethernet port6. In this configuration, the second NP4 processor would handle traffic acceleration and both XAUI links would be processing traffic.

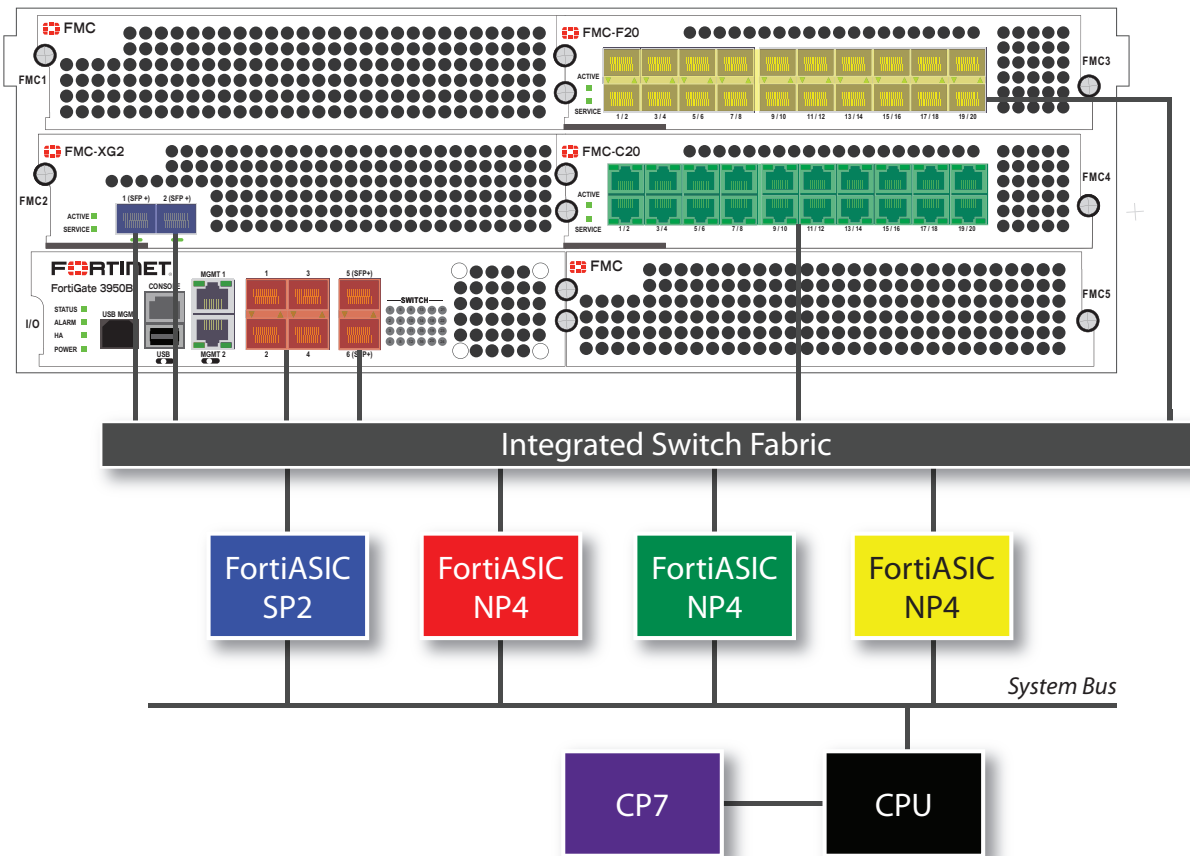
FortiGate-3950B and FortiGate-3951B

The FortiGate-3950B features one NP4 processor. The 1Gb SPF interfaces, port1, port2, port3, port4, and the 10Gb SPF+ interfaces, port5, port6, share connections to one NP4 processor. The FortiGate-3951B is similar to the FortiGate-3950B, except it trades one FMC slot for four FSM slots. The network interfaces available on each model are identical.

You can add additional FMC interface modules. The diagram below shows a FortiGate-3950B with three modules installed: an FMC-XG2, an FMC-F20, and an FMC-C20.

- The FMC-XG2 has one SP2 processor. The 10Gb SPF+ interfaces, port1 and port2, share connections to the processor.
- The FMC-F20 has one NP4 processor and the twenty 1Gb SPF interfaces, port1 through port20, share connections to the NP4 processor.

- The FMC-C20 has one NP4 processor and the twenty 10/100/1000 interfaces, port1 through port20, share connections to the NP4 processor.



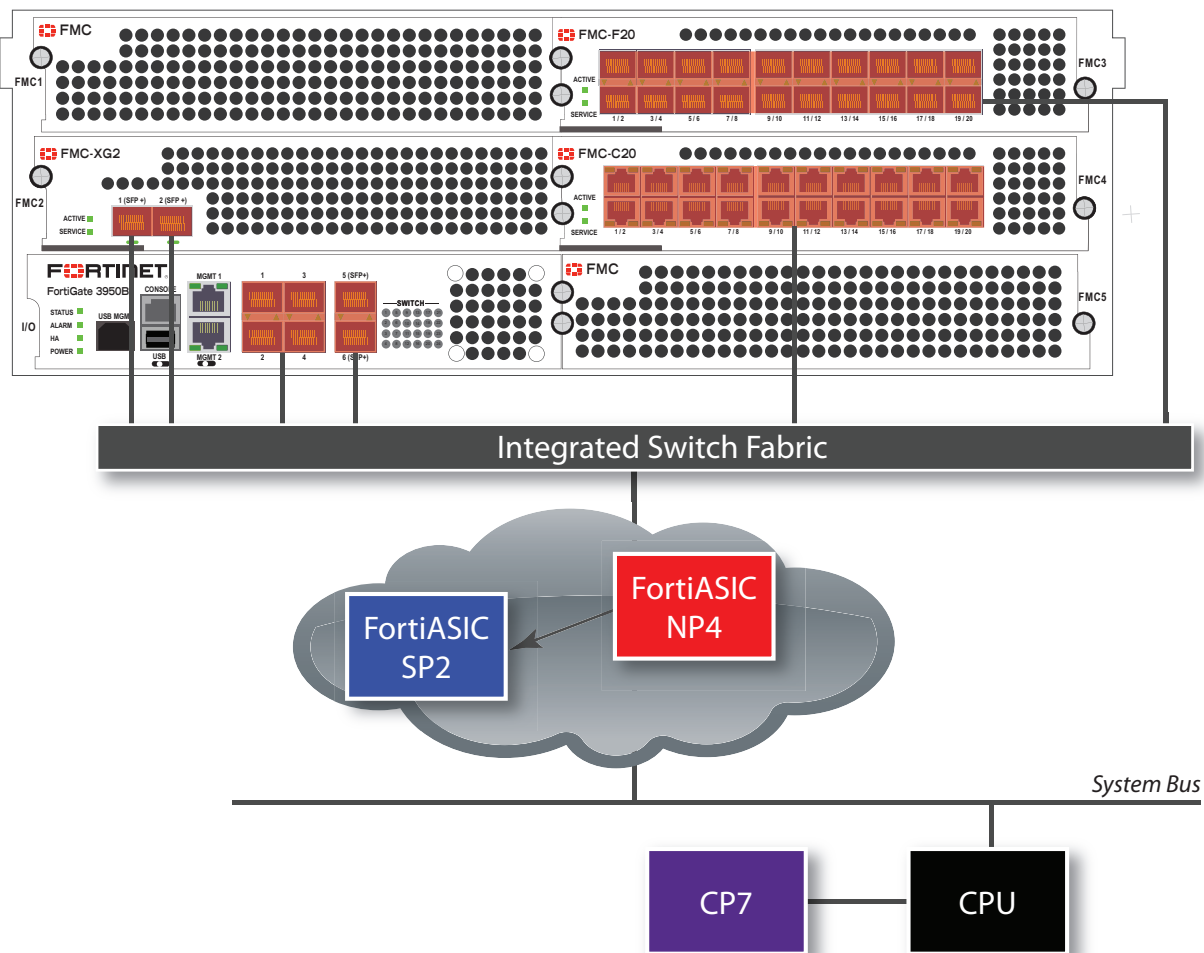
FortiGate-3950B and FortiGate-3951B — load balance mode

Adding one or more FMC-XG2 modules to your FortiGate-3950B allows you to enable load balance mode. This feature allows you increased flexibility in how you use the interfaces on the FortiGate unit. The FortiGate-3951B is similar to the FortiGate-3950B, except it trades one FMC slot for four FSM slots. The network interfaces available on each model are identical.

When enabled, traffic between any two interfaces (excluding management and console) is accelerated whether they are the six interfaces on the FortiGate-3950B itself, or on any installed FMC modules. Traffic is not limited to entering and leaving the FortiGate unit in specific interface groupings to benefit from NP4 and SP2 acceleration. You can use any pair of interfaces.

Security acceleration in this mode is limited, however. Only IPS scanning is accelerated in load balance mode.

The FortiGate-3950B in load balance mode



To enable this feature, issue this CLI command.

```
config system global
    set sp-load-balance enable
end
```

The FortiGate unit will then restart.

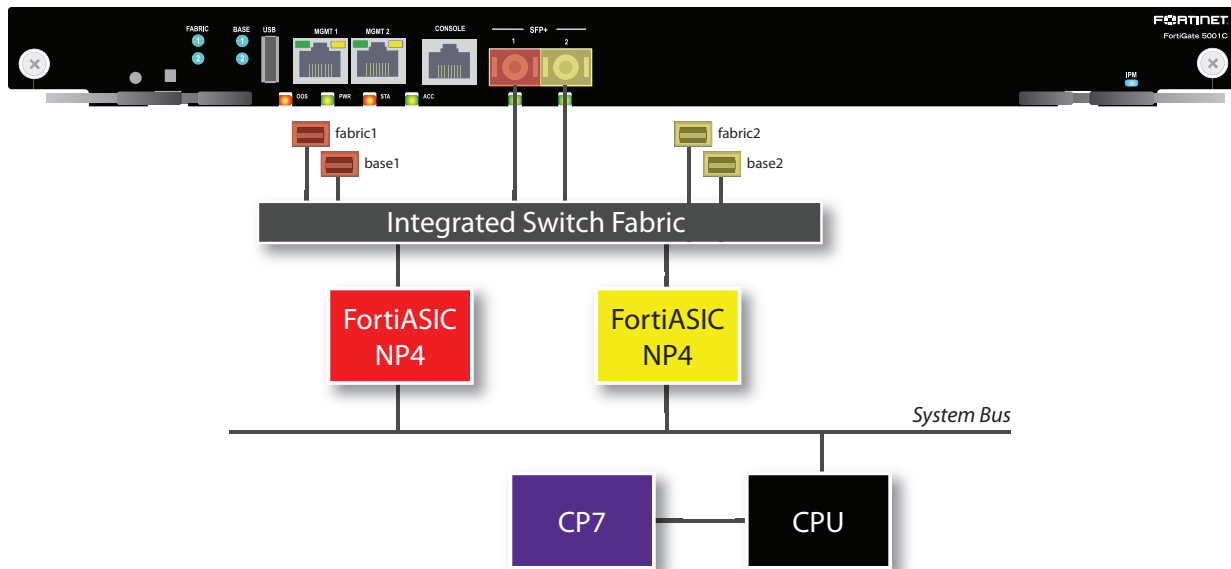
To return to the default mode, issue this CLI command.

```
config system global
    set sp-load-balance disable
end
```

FortiGate-5001C

The FortiGate-5001C board includes two NP4 processors connected to an integrated switch fabric:


- The port1, fabric1, and base1 interfaces are connected to one NP4 processor.
- The port2, fabric2, and base2 interfaces are connected to the other NP4 processor.



FortiGate-5001B

The FortiGate-5001B board includes two NP4 connected to an integrated switch fabric.


- The port1, port2, port3, port4, fabric1 and base1 interfaces are connected to one NP4 processor.
- The port5, port6, port7, port8, fabric2 and base2 interfaces are connected to the other NP4 processor.



Setting switch-mode mapping on the ADM-XD4

The ADM-XD4 SP has four 10Gb/s ports, but the NP4 processor it contains has only two 10Gb/s ports. The external ports you use are important to optimize the SP for your application.

ADM-XD4 mapping mode



Ports 1 and 3 share one NP4 processor and ports 2 and 4 share the other. Performance ports sharing the same NP4 processor is far better than when forcing network data to move between NP4 processors by using one port from each, for example ports 1 and 2 or ports 3 and 4.

Hardware acceleration get and diagnose commands

This section describes some `get` and `diagnose` commands you can use to display useful information about the NP6 processors sessions processed by NP6 processors.

get hardware npu np6

You can use the `get hardware npu np6` command to display information about the NP6 processors in your FortiGate and the sessions they are processing. This command contains a subset of the options available from the `diagnose npu np6` command. The command syntax is:

```
get hardware npu np6 {dce <np6-id> | ipsec-stats | port-list | session-stats <np6-id> |  
sse-stats <np6-id> | synproxy-stats}
```

`<np6-id>` identifies the NP6 processor. 0 is `np6_0`, 1 is `np6_1` and so on.

`dce` show NP6 non-zero sub-engine drop counters for the selected NP6.

`ipsec-stats` show overall NP6 IPsec offloading statistics.

`port-list` show the mapping between the FortiGate physical interfaces and NP6 processors.

`session-stats` show NP6 session offloading statistics counters for the selected NP6.

`sse-stats` show hardware session statistics counters.

`synproxy-stats` show overall NP6 synproxy statistics for TCP connections identified as being syn proxy DoS attacks.

diagnose npu np6

The `diagnose npu np6` command displays extensive information about NP6 processors and the sessions that they are processing. Some of the information displayed can be useful for understanding the NP6 configuration, seeing how sessions are being processed and diagnosing problems. Some of the commands may only be useful for Fortinet software developers. The command syntax is:

```
diagnose npu np6 {options}
```

The following options are available:

`fastpath {disable | enable} <np6-od>` enable or disable fastpath processing for a selected NP6.

`dce` shows NP6 non-zero sub-engine drop counters for the selected NP6.

`dce-all` show all subengine drop counters.

`anomaly-drop` show non-zero L3/L4 anomaly check drop counters.

`anomaly-drop-all` show all L3/L4 anomaly check drop counters.

`hrx-drop` show non-zero host interface drop counters.

`hrx-drop-all` show all host interface drop counters.

`session-stats` show session offloading statistics counters.

`session-stats-clear` clear session offloading statistics counters.

`sse-stats` show hardware session statistics counters.

`sse-stats-clear` show hardware session statistics counters.

`pdq` show packet buffer queue counters.

`xgmac-stats` show XGMAC MIBs counters.

`xgmac-stats-clear` clear XGMAC MIBS counters.

`port-list` show port list.

`ipsec-stats` show IPsec offloading statistics.

`ipsec-stats-clear` clear IPsec offloading statistics.

`EEPROM-read` read NP6 EEPROM.

`npu-feature` show NPU feature and status.

`register` show NP6 registers.

`fortilink` configure managed FortiSwitch.

`synproxy-stats` show synproxy statistics.

Using diagnose npu np6 npu-feature to verify enabled NP6 features

You can use the `diagnose npu np6 npu-feature` command to see what NP6 features are enabled and which are not. The following command output shows the normal default NP6 configuration for most FortiGates. In this output all features are enabled except low latency features and GRE offloading. Low latency is only available on the FortiGate-3700D and DX models and GRE offloading will become available in a future FortiOS release. The following output is from a FortiGate-1500D

```
diagnose npu np6 npu-feature
-----
```

	np_0	np_1
Fastpath	Enabled	Enabled
Low-latency-mode	Disabled	Disabled
Low-latency-cap	No	No
IPv4 firewall	Yes	Yes
IPv6 firewall	Yes	Yes
IPv4 IPsec	Yes	Yes
IPv6 IPsec	Yes	Yes
IPv4 tunnel	Yes	Yes
IPv6 tunnel	Yes	Yes
GRE tunnel	No	No
IPv4 Multicast	Yes	Yes
IPv6 Multicast	Yes	Yes
CAPWAP	Yes	Yes

If you use the following command to disable fastpath for np_0:

```
config system np6
  edit np6_0
    set fastpath disable
  end
```

The `npu-feature` command output show this configuration change:

```
diagnose npu np6 npu-feature
-----
np_0      np_1
-----
Fastpath      Disabled  Enabled
Low-latency-mode Disabled  Disabled
Low-latency-cap No        No
IPv4 firewall Yes       Yes
IPv6 firewall Yes       Yes
IPv4 IPsec    Yes       Yes
IPv6 IPsec    Yes       Yes
IPv4 tunnel   Yes       Yes
IPv6 tunnel   Yes       Yes
GRE tunnel    No        No
IPv4 Multicast Yes       Yes
IPv6 Multicast Yes       Yes
CAPWAP        Yes       Yes
```

Using the diagnose sys session/session6 list command

The `diagnose sys session list` and `diagnose sys session6 list` commands list all of the current IPv4 or IPv6 sessions being processed by the FortiGate. For each session the command output includes an `npu info` line that displays NPx offloading information for the session. If a session is not offloaded the command output includes a `no_ofld_reason` line that indicates why the session was not offloaded.

Displaying NP6 offloading information for a session

The `npu info` line of the `diagnose sys session list` command includes information about the offloaded session that indicates the type of processor and whether its IPsec or regular traffic:

- `offload=1/1` for NP1(FA1) sessions.
- `offload=2/2` for NP1(FA2) sessions.
- `offload=3/3` for NP2 sessions.
- `offload=4/4` for NP4 sessions.
- `offload=5/5` for XLR sessions.
- `offload=6/6` for NPlite/NP4lite sessions.
- `offload=7/7` for XLP sessions.
- `offload=8/8` for NP6 sessions.
- `flag 0x81` means regular traffic.
- `flag 0x82` means IPsec traffic.

Example offloaded IPv4 NP6 session

The following session output by the `diagnose sys session list` command shows an offloaded session. The information in the `npu info` line shows this is a regular session (`flag=0x81/0x81`) that is offloaded by an NP6 processor (`offload=8/8`).

```
diagnose sys session list
session info: proto=6 proto_state=01 duration=4599 expire=2753 timeout=3600 flags-
s=00000000 sockflag=00000000 sockport=0 av_idx=0 use=3
origin-shaper=
```

```

reply-shaper=
per_ip_shaper=
ha_id=0 policy_dir=0 tunnel=/ vlan_cos=0/255
state=log may_dirty npu none log-start
statistic(bytes/packets/allow_err): org=1549/20/1 reply=1090/15/1 tuples=2
speed(Bps/kbps): 0/0
origin->sink: org pre->post, reply pre->post dev=15->17/17->15
gwy=172.20.121.2/5.5.5.33
hook=post dir=org act=snat 5.5.5.33:60656->91.190.218.66:12350
(172.20.121.135:60656)
hook=pre dir=reply act=dnat 91.190.218.66:12350->172.20.121.135:60656
(5.5.5.33:60656)
pos/(before,after) 0/(0,0), 0/(0,0)
src_mac=98:90:96:af:89:b9
misc=0 policy_id=1 auth_info=0 chk_client_info=0 vd=0
serial=00058b9c tos=ff/ff app_list=0 app=0 url_cat=0
dd_type=0 dd_mode=0
npu_state=0x000c00
npu info: flag=0x81/0x81, offload=8/8, ips_offload=0/0, epid=140/138, ipid-
d=138/140, vlan=0x0000/0x0000
vlifid=138/140, vtag_in=0x0000/0x0000 in_npu=1/1, out_npu=1/1, fwd_en=0/0, qid=0/2

```

Example IPv4 session that is not offloaded

The following session, output by the diagnose sys session list command includes the `no_ofld_reason` line that indicates that the session was not offloaded because it is a local-in session.

```

session info: proto=6 proto_state=01 duration=19 expire=3597 timeout=3600
flags=00000000 sockflag=00000000 sockport=0 av_idx=0 use=3
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 policy_dir=0 tunnel=/ vlan_cos=8/8
state=local may_dirty
statistic(bytes/packets/allow_err): org=6338/15/1 reply=7129/12/1 tuples=2
speed(Bps/kbps): 680/5
origin->sink: org pre->in, reply out->post dev=15->50/50->15 gwy=5.5.5.5/0.0.0.0
hook=pre dir=org act=noop 5.5.5.33:60567->5.5.5.5:443(0.0.0.0:0)
hook=post dir=reply act=noop 5.5.5.5:443->5.5.5.33:60567(0.0.0.0:0)
pos/(before,after) 0/(0,0), 0/(0,0)
src_mac=98:90:96:af:89:b9
misc=0 policy_id=0 auth_info=0 chk_client_info=0 vd=0
serial=000645d8 tos=ff/ff app_list=0 app=0 url_cat=0
dd_type=0 dd_mode=0
npu_state=00000000
no_ofld_reason: local

```

Example IPv4 IPsec NP6 session

```

diagnose sys session list
session info: proto=6 proto_state=01 duration=34 expire=3565 timeout=3600 flags-
s=00000000 sockflag=00000000 sockport=0 av_idx=0 use=3
origin-shaper=

```

```

reply-shaper=
per_ip_shaper=
ha_id=0 policy_dir=0 tunnel=/p1-vdom2
state=re may_dirty npu
statistic(bytes/packets/allow_err): org=112/2/1 reply=112/2/1 tuples=2
origin->sink: org pre->post, reply pre->post dev=57->7/7->57 gwy-
y=10.1.100.11/11.11.11.1
hook=pre dir=org act=noop 172.16.200.55:35254->10.1.100.11:80(0.0.0.0:0)
hook=post dir=reply act=noop 10.1.100.11:80->172.16.200.55:35254(0.0.0.0:0)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=1 id_policy_id=0 auth_info=0 chk_client_info=0 vd=4
serial=00002d29 tos=ff/ff ips_view=0 app_list=0 app=0
dd_type=0 dd_mode=0
per_ip_bandwidth meter: addr=172.16.200.55, bps=260
npu_state=00000000
npu info: flag=0x81/0x82, offload=8/8, ips_offload=0/0, epid=1/3, ipid=3/1, vlan-
n=32779/0

```

Example IPv6 NP6 session

```

diagnose sys session6 list
session6 info: proto=6 proto_state=01 duration=2 expire=3597 timeout=3600 flags-
s=00000000 sockport=0 sockflag=0 use=3
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0
policy_dir=0 tunnel=/
state=may_dirty npu
statistic(bytes/packets/allow_err): org=152/2/0 reply=152/2/0 tuples=2
speed(Bps/kbps): 0/0
origin->sink: org pre->post, reply pre->post dev=13->14/14->13
hook=pre dir=org act=noop 2000:172:16:200::55:59145 ->2000:10:1:100::11:80(:::0)
hook=post dir=reply act=noop 2000:10:1:100::11:80 ->2000:172:16:200::55:59145
(:::0)
misc=0 policy_id=1 auth_info=0 chk_client_info=0 vd=0 serial=0000027a
npu_state=0x000c00
npu info: flag=0x81/0x81, offload=8/8, ips_offload=0/0, epid=137/136, ipid-
d=136/137, vlan=0/0

```

Example NAT46 NP6 session

```

diagnose sys session list
session info: proto=6 proto_state=01 duration=19 expire=3580 timeout=3600 flags-
s=00000000 sockflag=00000000 sockport=0 av_idx=0 use=3
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 policy_dir=0 tunnel=/
state=npu nlb
statistic(bytes/packets/allow_err): org=112/2/1 reply=112/2/1 tuples=2
speed(Bps/kbps): 0/0
origin->sink: org nataf->post, reply pre->org dev=52->14/14->52

```

```

gwy=0.0.0.0/10.1.100.1
hook=5 dir=org act=noop 10.1.100.1:21937->10.1.100.11:80(0.0.0.0:0)
hook=6 dir=reply act=noop 10.1.100.11:80->10.1.100.1:21937(0.0.0.0:0)
hook=pre dir=org act=noop 2000:172:16:200::55:33945 ->64:ff9b::a01:640b:80(:::0)
hook=post dir=reply act=noop 64:ff9b::a01:640b:80 ->2000:172:16:200::55:33945
(:::0)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=1 auth_info=0 chk_client_info=0 vd=0
serial=04051aae tos=ff/ff ips_view=0 app_list=0 app=0
dd_type=0 dd_mode=0
npu_state=00000000
npu info: flag=0x81/0x00, offload=0/8, ips_offload=0/0, epid=0/136, ipid=0/137,
vlan=0/0

```

Example NAT64 NP6 session

```

diagnose sys session6 list
session6 info: proto=6 proto_state=01 duration=36 expire=3563 timeout=3600 flags-
s=00000000 sockport=0 sockflag=0 use=3
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0
policy_dir=0 tunnel=/
state=may_dirty npu nlb
statistic(bytes/packets/allow_err): org=72/1/0 reply=152/2/0 tuples=2
speed(Bps/kbps): 0/0
orgin->sink: org pre->org, reply nataf->post dev=13->14/14->13
hook=pre dir=org act=noop 2000:172:16:200::55:33945 ->64:ff9b::a01:640b:80(:::0)
hook=post dir=reply act=noop 64:ff9b::a01:640b:80 ->2000:172:16:200::55:33945
(:::0)
hook=5 dir=org act=noop 10.1.100.1:21937->10.1.100.11:80(0.0.0.0:0)
hook=6 dir=reply act=noop 10.1.100.11:80->10.1.100.1:21937(0.0.0.0:0)
misc=0 policy_id=1 auth_info=0 chk_client_info=0 vd=0 serial=0000027b
npu_state=00000000
npu info: flag=0x00/0x81, offload=8/0, ips_offload=0/0, epid=137/0, ipid=136/0,
vlan=0/0

```

diagnose npu np6 session-stats <np6-id> (number of NP6 IPv4 and IPv6 sessions)

You can use the `diagnose npu np6 portlist` command to list the NP6-ids and the interfaces that each NP6 is connected to. The <np6-id> of np6_0 is 0, the <np6-id> of np6_1 is 1 and so on. The `diagnose npu np6 session-stats <np6-id>` command output includes the following headings:

- ins44 installed IPv4 sessions
- ins46 installed NAT46 sessions
- del4 deleted IPv4 and NAT46 sessions
- ins64 installed NAT64 sessions
- ins66 installed IPv6 sessions
- del6 deleted IPv6 and NAT64 sessions
- e is the error counter for each session type

```

diagnose npu np6 session-stats 0
qid    ins44      ins46      del4      ins64      ins66      del6
      ins44_e    ins46_e    del4_e    ins64_e    ins66_e    del6_e
-----
0      94        0          44        0          40         30
      0          0          0          0          0          0
1      84        0          32        0          30         28
      0          0          0          0          0          0
2      90        0          42        0          40         30
      0          0          0          0          0          0
3      86        0          32        0          24         27
      0          0          0          0          0          0
4      72        0          34        0          34         28
      0          0          0          0          0          0
5      86        0          30        0          28         32
      0          0          0          0          0          0
6      82        0          38        0          32         34
      0          0          0          0          0          0
7      86        0          30        0          30         30
      0          0          0          0          0          0
8      78        0          26        0          36         26
      0          0          0          0          0          0
9      86        0          34        0          32         32
      0          0          0          0          0          0
-----
Total 844        0          342        0          326        297
      0          0          0          0          0          0
-----

```

diagnose npu np6 ipsec-stats (NP6 IPsec statistics)

The command output includes IPv4, IPv6, and NAT46 IPsec information:

- `spi_ses4` is the IPv4 counter
- `spi_ses6` is the IPv6 counter
- `4to6_ses` is the NAT46 counter

```
diagnose npu np6 ipsec-stats
vif_start_oid      03ed      vif_end_oid      03fc
IPsec Virtual interface stats:
vif_get            000000000000      vif_get_expired  000000000000
vif_get_fail       000000000000      vif_get_invld    000000000000
vif_set            000000000000      vif_set_fail     000000000000
vif_clear          000000000000      vif_clear_fail   000000000000
np6_0:
sa_install         000000000000      sa_ins_fail      000000000000
sa_remove          000000000000      sa_del_fail      000000000000
4to6_ses_ins       000000000000      4to6_ses_ins_fail 000000000000
4to6_ses_del       000000000000      4to6_ses_del_fail 000000000000
spi_ses6_ins       000000000000      spi_ses6_ins_fail 000000000000
spi_ses6_del       000000000000      spi_ses6_del_fail 000000000000
spi_ses4_ins       000000000000      spi_ses4_ins_fail 000000000000
spi_ses4_del       000000000000      spi_ses4_del_fail 000000000000
sa_map_alloc_fail  000000000000      vif_alloc_fail   000000000000
sa_ins_null_adapter 000000000000      sa_del_null_adapter 000000000000
del_sa_mismatch    000000000000      ib_chk_null_adpt  000000000000
ib_chk_null_sa     000000000000      ob_chk_null_adpt  000000000000
ob_chk_null_sa     000000000000      rx_vif_miss      000000000000
rx_sa_miss         000000000000      rx_mark_miss     000000000000
waiting_ib_sa      000000000000      sa_mismatch       000000000000
msg_miss           000000000000
np6_1:
sa_install         000000000000      sa_ins_fail      000000000000
sa_remove          000000000000      sa_del_fail      000000000000
4to6_ses_ins       000000000000      4to6_ses_ins_fail 000000000000
4to6_ses_del       000000000000      4to6_ses_del_fail 000000000000
spi_ses6_ins       000000000000      spi_ses6_ins_fail 000000000000
spi_ses6_del       000000000000      spi_ses6_del_fail 000000000000
spi_ses4_ins       000000000000      spi_ses4_ins_fail 000000000000
spi_ses4_del       000000000000      spi_ses4_del_fail 000000000000
sa_map_alloc_fail  000000000000      vif_alloc_fail   000000000000
sa_ins_null_adapter 000000000000      sa_del_null_adapter 000000000000
del_sa_mismatch    000000000000      ib_chk_null_adpt  000000000000
ib_chk_null_sa     000000000000      ob_chk_null_adpt  000000000000
ob_chk_null_sa     000000000000      rx_vif_miss      000000000000
rx_sa_miss         000000000000      rx_mark_miss     000000000000
waiting_ib_sa      000000000000      sa_mismatch       000000000000
msg_miss           000000000000
```


diagnose sys mcast-session/session6 list (IPv4 and IPv6 multicast sessions)

This command lists all IPv4 or IPv6 multicast sessions. If a multicast session can be offloaded, the output includes the `offloadable` tag. If the multicast path can be offloaded one of the paths in the command output is tagged as `offloaded`.

The only way to determine the number of offloaded multicast sessions is to use the `diagnose sys mcast-session/session6 list` command and count the number of sessions with the `offload` tag.

```
diagnose sys mcast-session list
session info: id=3 vf=0 proto=17 172.16.200.55.51108->239.1.1.1.7878
used=2 path=11 duration=1 expire=178 indev=6 pkts=2 state:2cpu offloadable
npn-info in-pid=0 vifid=0 in-vtag=0 npuid=0 queue=0 tae=0

path: 2cpu policy=1, outdev=2
      out-vtag=0
path: 2cpu policy=1, outdev=3
      out-vtag=0
path: offloaded policy=1, outdev=7
      out-vtag=0
path: policy=1, outdev=8
      out-vtag=0
path: policy=1, outdev=9
      out-vtag=0
path: policy=1, outdev=10
      out-vtag=0
path: policy=1, outdev=11
      out-vtag=0
path: policy=1, outdev=12
      out-vtag=0
path: policy=1, outdev=13
      out-vtag=0
path: 2cpu policy=1, outdev=64
      out-vtag=0
path: 2cpu policy=1, outdev=68
      out-vtag=0
```

diagnose npu np6 sse-stats <np6-id> (number of NP6 sessions and dropped sessions)

This command displays the total number of inserted, deleted and purged sessions processed by a selected NP6 processor. The number of dropped sessions of each type can be determined by subtracting the number of successful sessions from the total number of sessions. For example, the total number of dropped insert sessions is `insert-total - insert-success`.

```
diagnose npu np6 sse-stats 0
```

Counters	SSE0	SSE1	Total
active	0	0	0
insert-total	25	0	0
insert-success	25	0	0
delete-total	25	0	0
delete-success	25	0	0
purge-total	0	0	0
purge-success	0	0	0
search-total	40956	38049	79005
search-hit	37714	29867	67581
pht-size	8421376	8421376	
oft-size	8355840	8355840	
oftfree	8355839	8355839	
PBA	3001		

diagnose npu np6 dce <np6-id> (number of dropped NP6 packets)

This command displays the number of dropped packets for the selected NP6 processor.

- `IHP1_PKTCHK` number of dropped IP packets
- `IPSEC0_ENGINB0` number of dropped IPSec
- `TPE_SHAPER` number of dropped traffic sharper packets

```
diag npu np6 dce 1
IHP1_PKTCHK :00000000000001833 [5b] IPSEC0_ENGINB0 :0000000000000003 [80]
TPE_SHAPER :00000000000000552 [94]
```

diagnose hardware deviceinfo nic <interface-name> (number of packets dropped by an interface)

This command displays a wide variety of statistics for FortiGate interfaces. The fields `Host Rx dropped` and `Host Tx dropped` display the number of received and transmitted packets that have been dropped.

```
diagnose hardware deviceinfo nic port2
```

```
...
```

```
===== Counters =====
```

```
Rx Pkts           :20482043
Rx Bytes          :31047522516
Tx Pkts           :19000495
Tx Bytes          :1393316953
Host Rx Pkts      :27324
Host Rx Bytes     :1602755
Host Rx dropped   :0
Host Tx Pkts      :8741
Host Tx Bytes     :5731300
Host Tx dropped   :0
sw_rx_pkts        :20482043
sw_rx_bytes       :31047522516
sw_tx_pkts        :19000495
sw_tx_bytes       :1393316953
sw_np_rx_pkts     :19000495
sw_np_rx_bytes    :1469318933
sw_np_tx_pkts     :20482042
sw_np_tx_bytes    :31129450620
```

diagnose npu np6 synproxy-stats (NP6 SYN-proxied sessions and unacknowledged SYNs)

This command displays information about NP6 syn-proxy sessions including the total number proxied sessions. As well the Number of attacks, no ACK from client shows the total number of acknowledged SYNs.

```
diagnose npu np6 synproxy-stats
DoS SYN-Proxy:
Number of proxied TCP connections : 39277346
Number of working proxied TCP connections : 182860
Number of retired TCP connections : 39094486
Number of attacks, no ACK from client : 208
```

Chapter 12 - High Availability

This FortiOS Handbook chapter contains the following sections:

[Solving the high availability problem](#) describes the high availability problem and introduces the FortiOS solutions described in this document (FGCP, VRRP, and standalone session synchronization).

[An introduction to the FGCP](#) introduces the FGCP clustering protocol and many of its features and terminology.

[FGCP configuration examples and troubleshooting](#) describes configuring HA clusters and contains HA clustering configuration examples.

[Virtual clusters](#) describes configuring HA virtual clusters and contains virtual clustering configuration examples.

[Full mesh HA](#) describes configuring FortiGate Full mesh HA and contains a full mesh HA configuration example.

[Operating clusters and virtual clusters](#) describes how to operate a cluster and includes detailed information about how various FortiGate systems operate differently in a cluster.

[HA and failover protection](#) describes in detail how FortiGate HA device failover, link failover, and session failover work.

[HA and load balancing](#) describes how FGCP HA active-active load balancing works and how to configure it.

[HA with FortiGate-VM and third-party products](#) describes how FortiGates interact with third-party products.

[VRRP high availability](#) describes FortiOS support of the Virtual Router Redundancy Protocol (VRRP) and its use for high availability.

[FortiGate Session Life Support Protocol \(FGSP\)](#) describes how to use the FGSP feature to support using external routers or load balancers to distribute or load balance sessions between two peer FortiGates.

What's new in FortiOS 6.0

The following list contains new high availability features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- Changes to get `system status ha` command output, see [Viewing cluster status from the CLI on page 1511](#)
- IPv6 support for VRRP, see [VRRP high availability on page 1595](#)
- [Session synchronization between FGCP clusters on page 1608](#)
- [GTP session synchronization: FGSP for FortiOS Carrier on page 1615](#)
- FGSP configuration synchronization changes, see [Synchronizing the configuration on page 1611](#)
- FGSP restoring a configuration is not synchronized, see [Backing up and restoring the configuration of an FGSP cluster on page 1613](#).
- New diagnose commands to find synchronization details, see [Determining what is causing a configuration synchronization problem on page 1545](#)
- HA IPv6 remote link monitoring, see [Configuring IPv6 remote IP monitoring on page 1558](#)

Solving the high availability problem

The basic high availability (HA) problem for TCP/IP networks and security gateways is keeping network traffic flowing. Uninterrupted traffic flow is a critical component for online systems and media because critical business processes quickly come to a halt when the network is down.

The security gateway is a crucial component of most networks since all traffic passes through it. A standalone network security gateway is a single point of failure that is vulnerable to any number of software or hardware problems that could compromise the device and bring all traffic on the network to a halt.

A common solution to the high availability problem is to eliminate the security gateway as single point of failure by introducing redundancy. With two or more redundant security gateways, if one fails, the remaining one or more gateways keep the traffic flowing. FortiOS provides six redundancy solutions: industry standard VRRP as well as five proprietary solutions: FortiGate Cluster Protocol (FGCP) high availability, FortiGate Session Life Support Protocol (FGSP) high availability, Session-Aware Load Balancing Clustering (SLBC), Enhanced Load Balanced Clustering (ELBC) and Content Clustering.

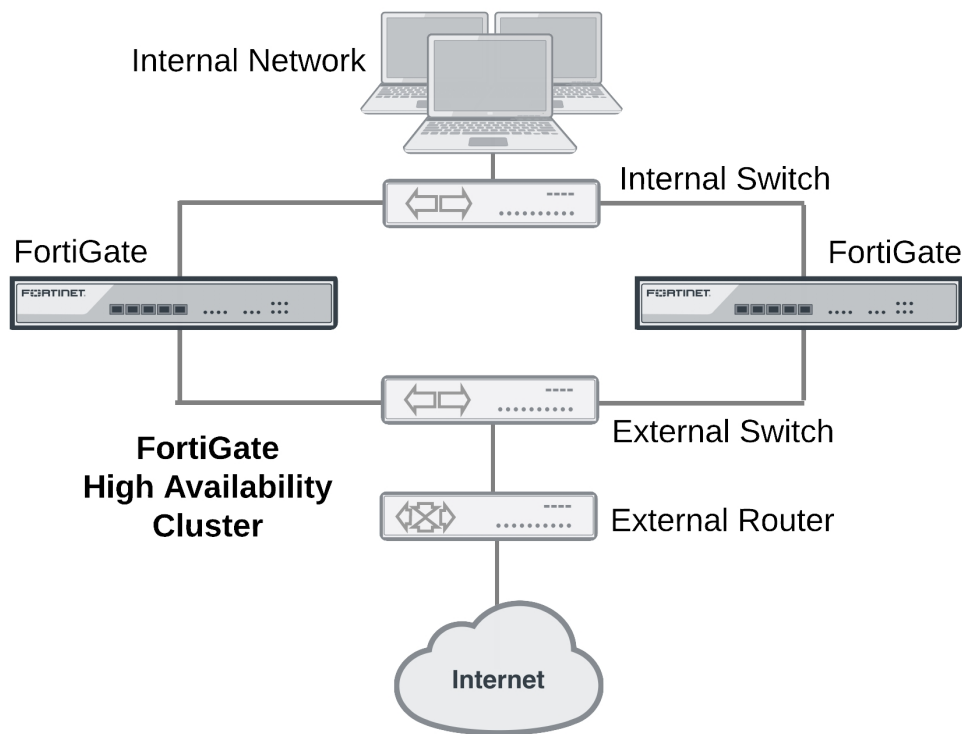


You can combine more than one high availability solution into a single configuration. A common reason for doing this could be to add VRRP to an FGCP or FGSP configuration.

A strong and flexible High availability solution is required for many mission-critical firewall and security profile applications. Each FortiOS high availability solution can be fine tuned to fit into many different network scenarios.

FortiGate Cluster Protocol (FGCP)

FGCP HA provides a solution for two key requirements of critical enterprise networking components: enhanced reliability and increased performance. Enhanced reliability is achieved through device failover protection, link failover protection and remote link failover protection. Also contributing to enhanced reliability is session failover protection for most IPv4 and IPv6 sessions including TCP, UDP, ICMP, IPsec VPN, and NAT sessions. Increased performance is achieved through active-active HA load balancing. Extended FGCP features include full mesh HA and virtual clustering. You can also fine tune the performance of the FGCP to change how a cluster forms and shares information among cluster units and how the cluster responds to failures.

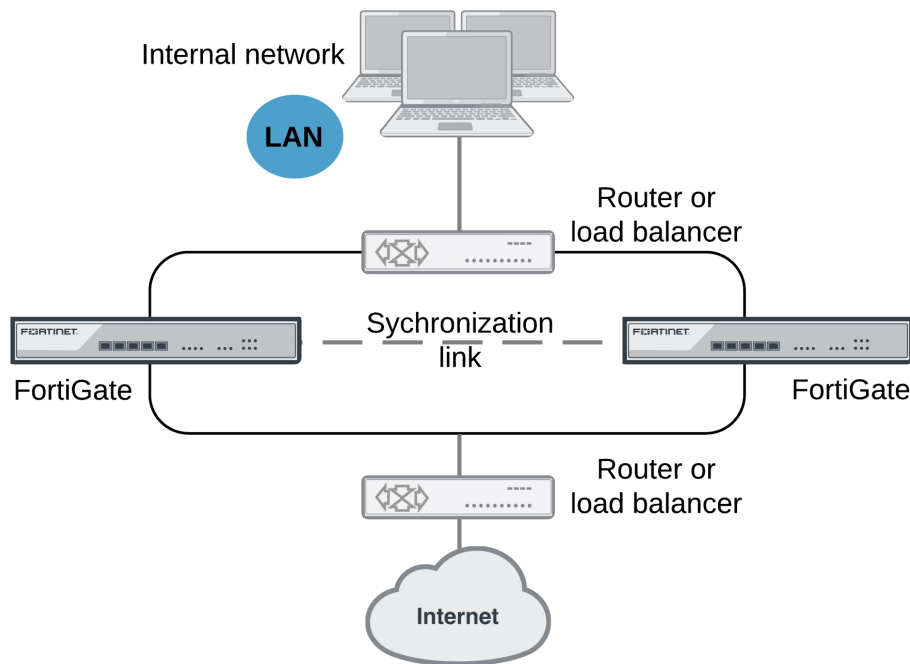


When configured onto your network an FGCP cluster appears to be a single FortiGate operating in NAT/Route or transparent mode and configuration synchronization allows you to configure a cluster in the same way as a standalone FortiGate. If a failover occurs, the cluster recovers quickly and automatically and also sends administrator notifications so that the problem that caused the failure can be corrected and any failed equipment restored.

The FGCP is compatible with most network environments and most networking equipment. While initial configuration is relatively quick and easy, a large number of tools and configuration options are available to fine tune the cluster for most situations.

FortiGate Session Life Support Protocol (FGSP)

In a network that already includes load balancing (either with load balancers or routers) for traffic redundancy, two identical FortiGates can be integrated into the load balancing configuration using the FortiGate Session Life Support Protocol (FGSP). The external load balancers or routers can distribute sessions among the FortiGates and the FGSP performs session synchronization of IPv4 and IPv6 TCP, SCTP, UDP, ICMP, expectation, and NAT sessions to keep the session tables of both FortiGates synchronized.



If one of the FortiGates fails, session failover occurs and active sessions fail over to the unit that is still operating. This failover occurs without any loss of data. As well, the external load balancers or routers detect the failover and re-distribute all sessions to the unit that is still operating.

Load balancing and session failover is done by external routers or load balancers and not by the FGSP. The FortiGates just perform session synchronization which allows session failover to occur without packet loss.

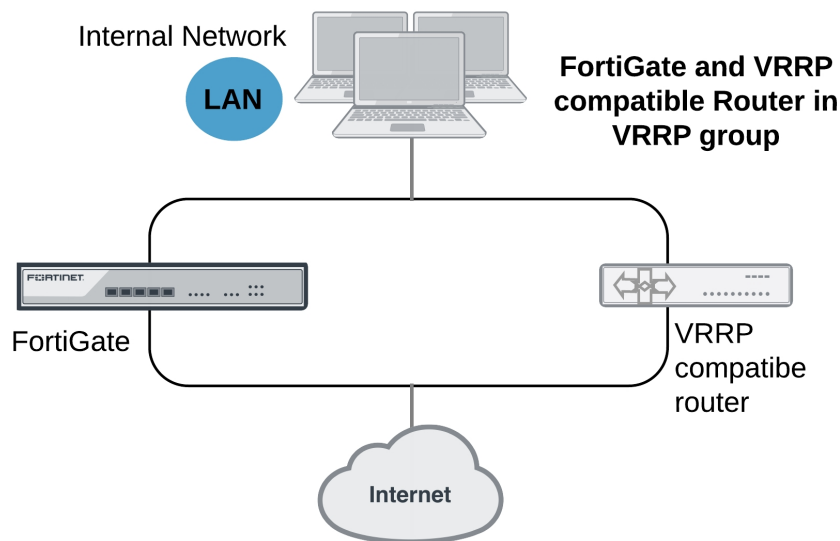
The FGSP also includes configuration synchronization, allowing you to make configuration changes once for both FortiGates instead of requiring duplicate configuration changes on each unit. However, settings that identify the FortiGate to the network, for example, interface IP addresses and BGP neighbor settings, are not synchronized so each FortiGate maintains its identity on the network. These settings must be configured separately for each FortiGate.



In previous versions of FortiOS the FGSP was called TCP session synchronization or standalone session synchronization. However, the FGSP has been expanded to include configuration synchronization and session synchronization of connectionless sessions, expectation sessions, and NAT sessions.

VRRP high availability

FortiGates can function as primary (master) or backup Virtual Router Redundancy Protocol (VRRP) routers and can be quickly and easily integrated into a network that has already deployed VRRP. A FortiGate can be integrated into a VRRP group with any third-party VRRP devices and VRRP can provide redundancy between multiple FortiGates. FortiOS supports VRRP version 2 and 3.



In a VRRP configuration, when a FortiGate operating as the primary router fails, a backup router takes its place and continues processing network traffic. If the backup router is a FortiGate, the network continues to benefit from FortiOS security features. If the backup router is simply a router, after a failure traffic will continue to flow, but FortiOS security features will be unavailable until the FortiGate is back on line. You can include different FortiGate models in the same VRRP group.

FortiOS supports IPv4 and IPv6 VRRP between two or more FortiGates and between FortiGates and third-party routers that support VRRP. Using VRRP, you can assign VRRP routers as primary or backup routers. The primary router processes traffic and the backup routers monitor the primary router and can begin forwarding traffic if the primary router fails. Similar to the FGCP, you can set up VRRP among multiple FortiGates to provide redundancy. You can also create a VRRP group with a FortiGate and any routers that support VRRP.

In a VRRP configuration that includes one FortiGate and one router, normally the FortiGate would be the primary router and all traffic would be processed by the FortiGate. If the FortiGate fails, all traffic switches to the router. Network connectivity is maintained even though FortiGate security features are unavailable until the FortiGate is back on line.

Session-Aware Load Balancing Clustering (SLBC)

Session-Aware Load Balancing Clusters consist of one or more FortiControllers acting as load balancers and two or more FortiGate-5000s and operating as workers all installed in one or two FortiGate-5000 series chassis.

SLBC clusters load balance TCP and UDP sessions. As a session-aware load balancer, the FortiController includes FortiASIC DP processors that maintain state information for all TCP and UDP sessions. The FortiASIC DP processors are capable of directing any TCP or UDP session to any worker installed in the same chassis. This session-awareness means that all TCP and UDP traffic being processed by a specific worker continues to be processed by the same worker. Session-awareness also means that more complex networking features such as network address translation (NAT), fragmented packets, complex UDP protocols, and complex protocols such as SIP that use pinholes, can be load balanced by the cluster.

For more information about SLBC see the *FortiController Session-Aware Load Balancing Guide*.



You cannot mix FGCP and SLBC clusters in the same FortiGate-5000 chassis.

Enhanced Load Balancing Clustering (ELBC)

ELBC uses FortiSwitch-5000 series load balancers to load balance traffic to FortiGate-5000 workers installed in a FortiGate-5000 chassis. ELBC enhances scalability, reliability, and availability of mission critical IP-based services, such as firewall, antivirus, web filtering, IPS, and so on. It also provides high availability by detecting worker failures and automatically redistributing traffic to the workers that are still operating.

ELBC applies a load balancing algorithm against the source and/or destination address of packets to generate a hash key value. Each worker has hash key values assigned to it. If the workers are running, then the traffic is forwarded to the worker assigned to the hash key. The hash key value generated by the algorithm, the hash keys accepted by the worker blades, and the blade the traffic is sent to are automatically calculated by the FortiSwitch.

For more information about ELBC see the *ELBC Configuration Guide*.



You cannot mix FGCP and ELBC clusters in the same FortiGate-5000 chassis.

Content clustering

A content cluster employs FortiSwitch-5203Bs or FortiController-5902Ds to load balance content sessions to FortiGate-5000 workers. FortiSwitch-5203B content clusters consist of one or more FortiSwitch-5203Bs and multiple FortiGate-5001Bs workers. FortiController-5902D content clusters consist of one or more FortiController-5902Ds and multiple FortiGate-5001B workers.

Operating as a FortiGate in content cluster mode, a primary FortiSwitch-5203B or FortiController-5902D performs routing, firewalling, stateful inspection, IPsec and SSL VPN encryption/decryption, and other FortiGate functions. The FortiSwitch-5203B includes NP4 processors and the FortiController-5902Ds includes NP6 processors and an integrated switch fabrics that provides fastpath acceleration by offloading communication sessions from the FortiGate CPU.

Using content cluster weighted load balancing, the FortiSwitch-5203Bs or FortiController-5902Ds distribute sessions that require content processing to the workers over the FortiGate-5000 chassis fabric backplane. Content processing sessions include proxy and flow-based security profile functions such as virus scanning, intrusion protection, application control, IPS, web filtering, email filtering, and VoIP. Load balancing is offloaded to the NP4 or NP6 processors resulting in improved load balancing performance. In some networks, the NP4 or NP6 processors also allow you to configure the efficiently load balance TCP and UDP sessions.

Content cluster mode is similar to active-active HA where the FortiSwitch-5203B or FortiController-5902D operates as the primary unit and load balances security profile sessions to the workers installed in the chassis using weighted load balancing. In this configuration, the HA mode is active-active, the HA load balancing schedule is weight-round-robin and load-balance-all is disabled. You can adjust the HA weighted load balancing weights to change how sessions are load balanced.

You can add a second FortiSwitch-5203B or FortiController-5902D to a content cluster as a backup. The primary FortiSwitch-5203B or FortiController-5902D can load balance sessions to the backup FortiSwitch-5203B or

FortiController-5902D as well as the workers. You can control how many sessions are processed by the backup FortiSwitch-5203B or FortiController-5902D by configuring the HA load balancing weights. You can also configure the content cluster to operate the backup FortiSwitch-5203B or FortiController-5902D in standby mode. In this mode the backup FortiSwitch-5203B or FortiController-5902D does not process any sessions but is just there to take over content clustering if the primary unit fails.

Once the content cluster has been established and all FortiControllers and workers have joined the cluster, you can configure the cluster from the FortiSwitch-5203B or FortiController-5902D GUI or CLI. All configuration changes made to the primary unit are automatically synchronized to all cluster units.

FortiSwitch-5203B or FortiController-5902D firmware upgrades are done from the primary FortiSwitch-5203B or FortiController-5902D GUI or CLI. Worker firmware upgrades are done from the FortiSwitch-5203B or FortiController-5902D CLI where a single firmware image is uploaded once and synchronized to all of the workers.

An introduction to the FGCP

A FortiGate HA cluster consists of two to four FortiGates configured for HA operation. Each FortiGate in a cluster is called a cluster unit. All cluster units must be the same FortiGate model with the same FortiOS firmware build installed. All cluster units must also have the same hardware configuration (for example, the same number of hard disks and so on) and be running in the same operating mode (NAT/Route mode or transparent mode).



You can create an FGCP cluster of up to four FortiGates.

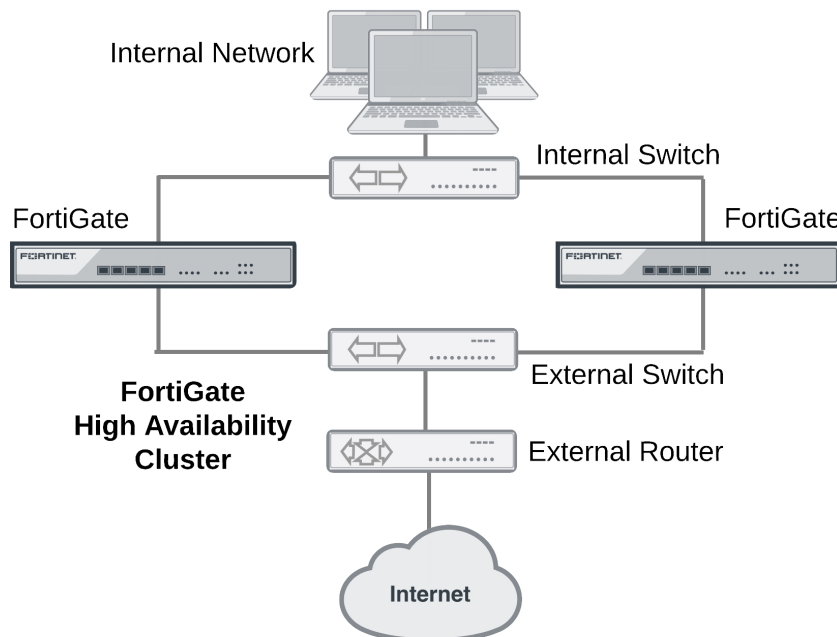
In addition the cluster units must be able to communicate with each other through their heartbeat interfaces. This heartbeat communication is required for the cluster to be created and to continue operating. Without it, the cluster acts like a collection of standalone FortiGates.

On startup, after configuring the cluster units with the same HA configuration and connecting their heartbeat interfaces, the cluster units use the FortiGate Clustering Protocol (FGCP) to find other FortiGates configured for HA operation and to negotiate to create a cluster. During cluster operation, the FGCP shares communication and synchronization information among the cluster units over the heartbeat interface link. This communication and synchronization is called the FGCP heartbeat or the HA heartbeat. Often, this is shortened to just heartbeat.

The cluster uses the FGCP to select the primary unit, and to provide device, link and session failover. The FGCP also manages the two HA modes; active-passive (failover HA) and active-active (load balancing HA).

About the FGCP

FortiGate HA is implemented by configuring two or more FortiGates to operate as an HA cluster. To the network, the HA cluster appears to function as a single FortiGate, processing network traffic and providing normal security services such as firewalling, security profile services, Security Fabric services, and VPN services.

HA cluster installed between an internal network and the internet

Inside the cluster the individual FortiGates are called cluster units. These cluster units share state and configuration information. If one cluster unit fails, the other units in the cluster automatically replace that unit, taking over the work that the failed unit was doing. After the failure, the cluster continues to process network traffic and provide normal FortiGate services with virtually no interruption.

Every FortiGate cluster contains one primary unit (also called the master unit) and one or more subordinate units (also called slave or backup units). The primary unit controls how the cluster operates. The role that the subordinate units play depends on the mode in which the cluster operates: (Active-Passive (AP) or Active-Active (AA).

The ability of an HA cluster to continue providing firewall services after a failure is called failover. FGCP failover means that your network does not have to rely on one FortiGate to continue functioning. You can install additional units and form an HA cluster.

A second HA feature, called active-active load balancing, can be used to increase performance. An active-active cluster of FortiGates can increase overall network performance by sharing the load of processing network traffic and providing security services. The cluster appears to your network to be a single device, adding increased performance without changing your network configuration.

Virtual clustering extends HA features to provide failover protection and load balancing for Virtual Domains (VDOMs). See [Virtual clusters on page 1463](#).

FortiGate models that support redundant interfaces can be configured to support full mesh HA. Full mesh HA is a method of reducing the number of single points of failure on a network that includes an HA cluster. For details about full mesh HA, see [Full mesh HA on page 1473](#).

FGCP failover protection

The FGCP provides IP/MAC takeover failover protection by assigning virtual MAC addresses to the primary cluster unit and then sending gratuitous ARP packets from the primary unit interfaces to reprogram the network.

Failover times can be less than a second under optimal conditions. You can fine tune failover performance for your network by adjusting cluster status checking timers, routing table update timers, and wait timers.

An HA cluster fails over if the primary unit fails (a device failure) or experiences a link failure. The cluster can detect link failures for connections to the primary unit using port monitoring and for connections between downstream network components using remote IP monitoring. To compensate for a link failover, the cluster maintains active links to keep traffic flowing between high-priority networks. Port and remote IP monitoring can be fine tuned without disrupting cluster operation.

Session failover

FGCP session failover maintains TCP, SIP and IPsec VPN sessions after a failure. You can also configure session failover to maintain UDP and ICMP sessions. Session failover does not failover SSL VPN sessions. Session failover may not be required for all networks because many TCP/IP, UDP, and ICMP protocols can resume sessions on their own. Supporting session failover adds extra overhead to cluster operations and can be disabled to improve cluster performance if it is not required.

Load balancing

Active-active HA load balances resource-intensive security profile features such as virus scanning, web filtering, intrusion protection, application control, email filtering and data leak prevention operations among all cluster units to provide better performance than a standalone FortiGate. If network traffic consists of mainly TCP sessions, the FGCP can also load balance all TCP sessions to improve TCP performance in some network configurations. On some FortiGate models you can also load balance UDP sessions. NP4 and NP6 offloading can accelerate HA load balancing (especially TCP and UDP load balancing). HA load balancing schedules can be adjusted to optimize performance for the traffic mix on your network. Weighted load balancing can be used to control the relative amount of sessions processed by each cluster unit.

Virtual clustering

Virtual clustering is an extension of the FGCP for a cluster of 2 FortiGates operating with multiple VDOMS enabled. Not only does virtual clustering provide failover protection for a multiple VDOM configuration, but a virtual cluster can load balance traffic between the cluster units. Load balancing with virtual clustering is quite efficient and load balances all traffic. It is possible to fine tune virtual clustering load balancing in real time to actively optimize load sharing between the cluster units without affecting the smooth operation of the cluster.

Full mesh HA

High availability improves the reliability of a network by replacing a single point of failure (a single FortiGate) with a cluster that can maintain network traffic if one of the cluster units fails. However, in a normal FGCP cluster, single points of failure remain. Full mesh HA removes these single points of failure by allowing you to connect redundant switches to each cluster interface. Full mesh HA is achieved by configuring 802.3ad aggregate or redundant interfaces on the FortiGate and connecting redundant switches to these interfaces. Configuration is a relatively simple extension of the normal aggregate interface and HA configurations.

Cluster management

FortiOS HA provides a wide range of cluster management features:

- Automatic continuous configuration synchronization. You can get a cluster up and running almost as quickly as a standalone FortiGate by performing a few basic steps to configure HA settings and minimal network settings on each cluster unit. When the cluster is operating you can configure FortiGate features such as firewalling, content inspection, and VPN in the same way as for a standalone FortiGate. All configuration changes (even complex changes such as switching to multiple VDOM mode or from NAT/Route to transparent mode) are synchronized among all cluster units.
- Firmware upgrades/downgrades. Upgrading or downgrading cluster firmware is similar to upgrading or downgrading standalone FortiGate firmware. The Firmware is uploaded once to the primary unit and the cluster automatically upgrades or downgrades all cluster units in one operation with minimal or no service interruption.
- Individual cluster unit management. In some cases you may want to manage individual cluster units. You can do so from cluster CLI by navigating to each cluster unit. You can also use the reserved management interface feature to give each cluster unit its own IP address and default route. You can use the reserved management interfaces and IP addresses to connect to the GUI and CLI of each cluster unit and configure an SNMP server to poll each cluster unit.
- Removing and adding cluster units. In one simple step any unit (even the primary unit) can be removed from a cluster and given a new IP address. The cluster keeps operating as it was; the transition happening without interrupting cluster operation. A new unit can also be added to an operating cluster without disrupting network traffic. All you have to do is connect the new unit and change its HA configuration to match the cluster's. The cluster automatically finds and adds the unit and synchronizes its configuration with the cluster.
- Debug and diagnose commands. An extensive range of debug and diagnose commands can be used to report on HA operation and find and fix problems.
- Logging and reporting. All cluster units can be configured to record all log messages. These message can be stored on the individual cluster units or sent to a FortiAnalyzer unit. You can view all cluster unit log messages by logging into any cluster unit.
- FortiManager support. FortiManager understands FortiOS HA and automatically recognizes when you add a FortiOS cluster to the FortiManager configuration.

The FGCP uses a combination of incremental and periodic synchronization to make sure that the configuration of all cluster units is synchronized to that of the primary unit. This means that in most cases you only have to make a configuration change once to have it synchronized to all cluster units.

Synchronizing the configuration (and settings that are not synchronized)

The FGCP uses a combination of incremental and periodic synchronization to make sure that the configuration of all cluster units is synchronized to that of the primary unit. This means that in most cases you only have to make a configuration change once to have it synchronized to all cluster units. This includes special configuration settings that include extra information (for example, 3rd party certificates, replacement message text files and graphics and so on).

Some configuration settings are not synchronized to support some aspects of FortiGate operation. The following settings are not synchronized among cluster units:

- The FortiGate host name. Allows you to identify cluster units.
- The GUI Dashboard configuration. After a failover you may have to re-configure dashboard widgets.
- HA override.
- HA device priority.
- Virtual cluster 1 and Virtual cluster 2 device priorities.
- The HA priority (`ha-priority`) setting for a ping server or dead gateway detection configuration.
- The system interface settings of the FortiGate interface that becomes the HA reserved management interface.

- The default route for the reserved management interface, set using the `ha-mgmt-interface-gateway` option of the `config system ha` command.
- The dynamic weighted load balancing thresholds and high and low watermarks.
- OSPF `summary-addresses` settings.

In addition licenses are not synchronized since each FortiGate must be licensed separately. This includes FortiCloud activation and FortiClient licensing, and entering a license key if you purchased more than 10 Virtual Domains (VDOMS).

Preparing the FortiGates before setting up an FGCP cluster

Before creating an FGCP cluster you should complete the following setup on each FortiGate.

DHCP and PPPoE

Make sure your FortiGate interfaces are configured with static IP addresses. If any interface gets its address using DHCP or PPPoE you should temporarily switch it to a static address and enable DHCP or PPPoE after the cluster has been established.

Firmware version

Make sure the FortiGates are running the same FortiOS firmware version.

About HA and licensing

All of the FortiGates in a cluster must have the same level of licensing. This includes licensing for FortiCare Support, IPS, AntiVirus, Web Filtering, Mobile Malware, FortiClient, FortiCloud, and additional virtual domains (VDOMS).

If one of the FortiGates in a cluster has a lower level of licensing than other FortiGates in the cluster, then all of the FortiGates in the cluster will revert to that lower licensing level. For example, if you only purchase FortiGuard Web Filtering for one of the FortiGates in a cluster, when the cluster is operating, none of the cluster units will support FortiGuard Web Filtering.

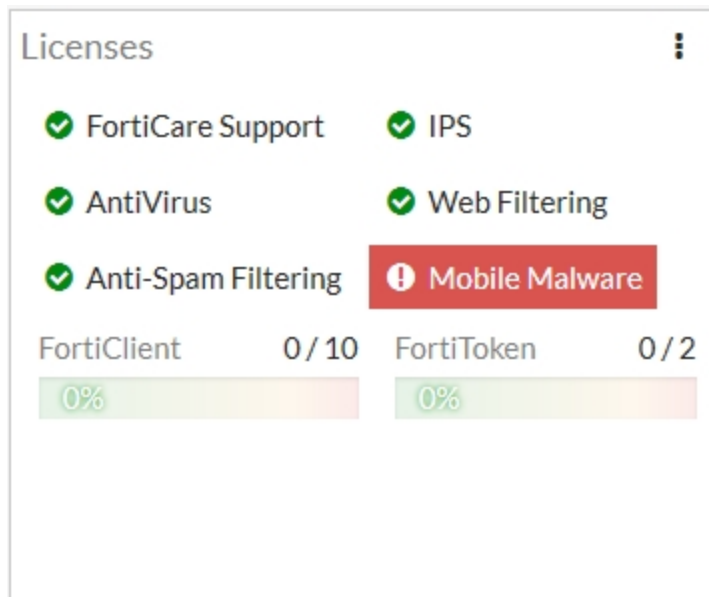
An exception is FortiToken licensing. FortiToken activations are completed one FortiGate unit and synchronized to all of the FortiGates in the cluster.

FortiOS Carrier license

If the FortiGates in the cluster will be running FortiOS Carrier, apply the FortiOS Carrier license before configuring the cluster (and before applying other licenses). Applying the FortiOS Carrier license sets the configuration to factory defaults, requiring you to repeat steps performed before applying the license. All FortiGates in the cluster must be licensed for FortiOS Carrier.

Support contracts and FortiGuard, FortiCloud, FortiClient, and VDOM licensing

Register and apply these licenses to each FortiGate. This includes FortiCloud activation and FortiClient licensing, and entering a license key if you purchased more than 10 Virtual Domains (VDOMS). All FortiGates in the cluster must have the same level of licensing for FortiGuard, FortiCloud, FortiClient and VDOMs.



FortiToken licenses

You only need one set of FortiToken licenses for the HA cluster and you only need to activate each token once. Normally you would activate your tokens on the primary unit and this configuration and the seed information will be synchronized to all cluster members so all tokens will then be activated for all cluster members.

If you have added FortiToken licenses and activated FortiTokens on a standalone FortiGate unit before configuring HA, the licenses and the FortiToken activations will usually be synchronized to all cluster units after forming a cluster. To make sure this goes smoothly you can make sure the FortiGate that you have added the licenses to becomes the primary unit when setting up the cluster as described in [How to set up FGCP clustering \(recommended steps\)](#) on page 1412.

Certificates

You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate.

Configuring FortiGates for FGCP HA operation

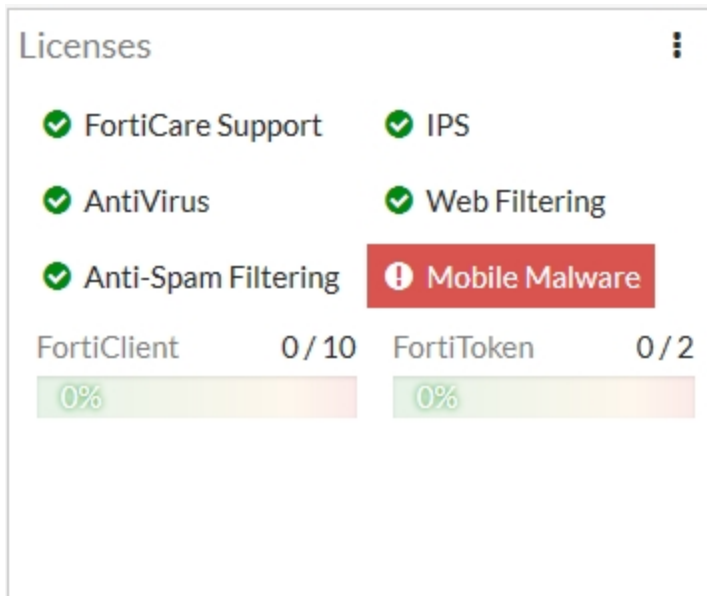
Each FortiGate in the cluster must have the same HA configuration. Once the cluster is connected, you can configure it in the same way as you would configure a standalone FortiGate. The following example sets the HA mode to active-passive and the HA password to HA_pass.



Make sure your FortiGate interfaces are configured with static IP addresses. If any interface gets its address using DHCP or PPPoE you should temporarily switch it to a static address and enable DHCP or PPPoE after the cluster has been established.

Make sure both FortiGates are running the same FortiOS firmware version. Register and apply licenses to both FortiGates before adding them to the cluster. This includes licensing for FortiCare Support, IPS, AntiVirus, Web Filtering, Mobile Malware, FortiClient, FortiCloud, and additional virtual domains (VDOMs). All FortiGates in the

cluster must have the same level of licensing for FortiGuard, FortiCloud, FortiClient, and VDOMs. FortiToken licenses can be added at any time because they are synchronized to all cluster members.



You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed, third-party certificates are synchronized to the backup FortiGate.

To configure a FortiGate for HA operation - GUI

1. Power on the FortiGate to be configured.
2. Log into the GUI.
3. Locate the System Information Dashboard widget. Click on the System Information dashboard widget and select **Configure settings in System > Settings**.
4. Enter a new Host Name for this FortiGate.
Changing the host name makes it easier to identify individual cluster units when the cluster is operating.
5. Go to **System > HA** and change the following settings:

Mode	Active-Passive
Group Name	Example_cluster
Password	HA_pass
The password must be the same for all FortiGates in the cluster.	

You can accept the default configuration for the remaining HA options and change them later, once the cluster is operating.

6. Select **OK**.
The FortiGate negotiates to establish an HA cluster. When you select **OK** you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all ARP table entries). You may be able to delete the ARP table

of your management PC from a command prompt using a command similar to `arp -d`.

7. Power off the FortiGate.
8. Repeat this procedure for all of the FortiGates in the cluster.
Once all of the units are configured, continue by connecting the FortiGate HA cluster below.

To configure a FortiGate for HA operation - CLI

1. Power on the FortiGate to be configured.
2. Log into the CLI.
3. Enter the following command to change the FortiGate host name.

```
config system global
  set hostname Example1_host
end
```

Changing the host name makes it easier to identify individual cluster units when the cluster is operating.

4. Enter the following command to enable HA:

```
config system ha
  set mode active-passive
  set group-name Example_cluster
  set password HA_pass
end
```

You can accept the default configuration for the remaining HA options and change them later, once the cluster is operating.

The FortiGate negotiates to establish an HA cluster. You may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate interfaces. To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

5. Power off the FortiGate.
6. Repeat this procedure for all of the FortiGates in the cluster.
Once all of the units are configured, continue with connecting the FortiGate HA cluster.

Connecting a FortiGate HA cluster

Use the following procedure to connect a cluster. Connect the cluster units to each other and to your network. You must connect all matching interfaces in the cluster to the same switch, then connect these interfaces to their networks using the same switch.

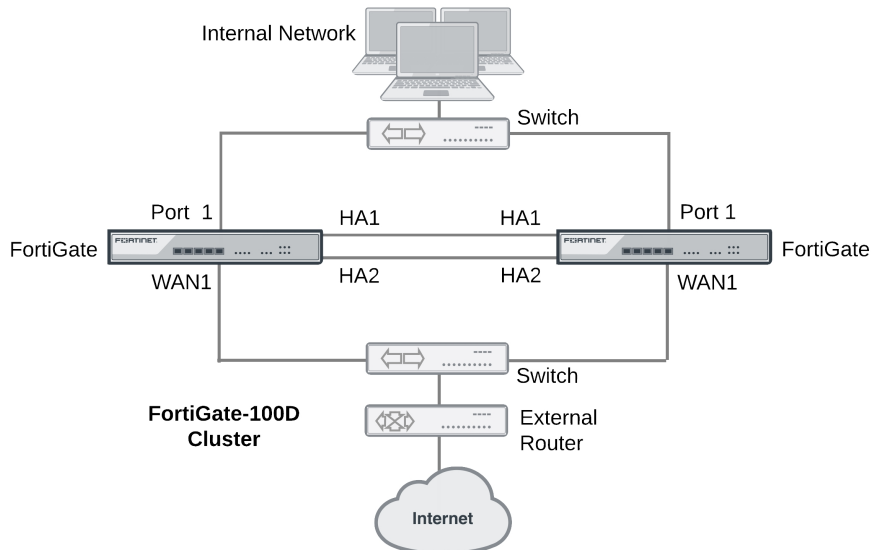
Although you can use hubs, Fortinet recommends using switches for all cluster connections for the best performance.

Connecting an HA cluster to your network temporarily interrupts communications on the network because new physical connections are being made to route traffic through the cluster. Also, starting the cluster interrupts network traffic until the individual cluster units are functioning and the cluster completes negotiation. Cluster negotiation is automatic and normally takes just a few seconds. During system startup and negotiation all network traffic is dropped.

This section describes how to connect the cluster shown below, which consists of two FortiGate-100D units to be connected between the internet and a head office internal network. The wan1 interfaces of the FortiGate connect

the cluster to the internet and the internal interfaces connect the cluster to the internal network. The ha1 and ha2 interfaces are used for redundant HA heartbeat links.

Example cluster connections



To connect a FortiGate HA cluster

1. Connect the WAN1 interfaces of each cluster unit to a switch connected to the internet.
2. Connect the Port1 interfaces of each cluster unit to a switch connected to the internal network.
3. Connect the HA1 interfaces of the cluster units together. You can use a crossover Ethernet cable or a regular Ethernet cable. (You can also connect the interfaces using Ethernet cables and a switch.)
4. Connect the HA2 interfaces of the cluster units together. You can use a crossover Ethernet cable or a regular Ethernet cable. (You can also connect the interfaces using Ethernet cables and a switch.)
5. Power on both of the FortiGates.

As the cluster units start, they negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally just takes a few seconds.

At least one heartbeat interface should be connected together for the cluster to operate.

Do not use a switch port for the HA heartbeat traffic. This configuration is not supported.

You could use one switch to connect all four heartbeat interfaces. However, this is not recommended because if the switch fails both heartbeat interfaces will become disconnected.

6. You can now configure the cluster as if it is a single FortiGate.

Verifying the cluster status from the HA Status dashboard widget

The HA Status dashboard widget shows the mode and group names of the cluster, the status of the cluster units and their host names, the cluster uptime and the last time the cluster state changed. A state change can indicate the cluster first forming or one of the cluster units changing its role in the cluster.

The HA Status Dashboard widget also shows if the cluster units are synchronized. Mouse over each FortiGate in the cluster to verify that they both have the same checksum.

From the HA Status widget you can also select **Show HA Historical Events** to see the most recent HA system status messages.

HA Historical Events		
#	Date/Time	Event
1	10-05 14:41	FGT6HD3916800525 is elected as the cluster master of 2 members
2	10-05 14:41	new member 'FGT6HD3916801195' joins the cluster
3	10-05 14:41	hbdev port4 link status changed: 0->1
4	10-05 14:41	port port4 link status changed: 0->1
5	10-05 14:41	hbdev port4 link status changed: 1->0
6	10-05 14:41	port port4 link status changed: 1->0
7	10-05 14:41	hbdev port4 link status changed: 0->1
8	10-05 14:41	hbdev port3 link status changed: 0->1
9	10-05 14:41	port port4 link status changed: 0->1
10	10-05 14:41	port port3 link status changed: 0->1
11	10-05 14:40	hbdev port4 link status changed: 1->0
12	10-05 14:40	hbdev port3 link status changed: 1->0
13	10-05 14:40	port port4 link status changed: 1->0

Active-passive and active-active HA

The first decision to make when configuring FortiGate HA is whether to choose active-passive or active-active HA mode. To configure the HA mode, go to **System > HA** and set Mode to **Active-Passive** or **Active-Active**.

From the CLI enter the following command to set the HA mode to active-passive:

```
config system ha
  set mode a-p
end
```

To form a cluster, all cluster units must be set to the same mode. You can also change the mode after the cluster is up and running. Changing the mode of a functioning cluster causes a slight delay while the cluster renegotiates to operate in the new mode and possibly select a new primary unit.

Active-passive HA (failover protection)

An active-passive (A-P) HA cluster provides hot standby failover protection. An active-passive cluster consists of a primary unit that processes communication sessions, and one or more subordinate units. The subordinate units are connected to the network and to the primary unit but do not process communication sessions. Instead, the subordinate units run in a standby state. In this standby state, the configuration of the subordinate units is synchronized with the configuration of the primary unit and the subordinate units monitor the status of the primary unit.

Active-passive HA provides transparent device failover among cluster units. If a cluster unit fails, another immediately take its place.

Active-passive HA also provides transparent link failover among cluster units. If a cluster unit interface fails or is disconnected, this cluster unit updates the link state database and the cluster negotiates and may select a new primary unit.

If session failover (also called session pickup) is enabled, active-passive HA provides session failover for some communication sessions.

The following example shows how to configure a FortiGate for active-passive HA operation. You would enter the exact same commands on every FortiGate in the cluster.

```
config system ha
  set mode a-p
  set group-name myname
  set password HApass
end
```

Active-active HA (load balancing and failover protection)

By default, active-active HA load balancing distributes proxy-based security profile processing to all cluster units. Proxy-based security profile processing is CPU and memory-intensive, so FGCP load balancing may result in higher throughput because resource-intensive processing is distributed among all cluster units.

Normally, sessions accepted by policies that don't include security profiles are not load balanced and are processed by the primary unit. You can configure active-active HA to load balance additional sessions.

An active-active HA cluster consists of a primary unit that receives all communication sessions and load balances them among the primary unit and all of the subordinate units. In an active-active cluster the subordinate units are also considered active since they also process content processing sessions. In all other ways active-active HA operates the same as active-passive HA.

The following example shows how to configure a FortiGate for active-active HA operation. You would enter the exact same commands on every FortiGate in the cluster.

```
config system ha
  set mode a-a
  set group-name myname
  set password HApass
end
```

Identifying the cluster and cluster units

You can use the cluster group name, group id, and password to identify a cluster and distinguish one cluster from another. If you have more than one cluster on the same network, each cluster must have a different group name, group id, and password.

Group name

Use the group name to identify the cluster. The maximum length of the group name is 32 characters. The group name must be the same for all cluster units before the cluster units can form a cluster. After a cluster is operating, you can change the group name. The group name change is synchronized to all cluster units. The group name appears on the HA Status dashboard widget.

To add or change the group name from the GUI go to **System > HA** and change the **Group name**.

Enter the following CLI command to change the group name to Cluster_name:

```
config system ha
    set group-name Cluster_name
end
```

Password

Use the password to identify the cluster. You should always change the password when configuring a cluster. The password must be the same for all FortiGates before they can form a cluster. When the cluster is operating you can change the password, if required. Two clusters on the same network cannot have the same password.

To change the password from the GUI go to **System > HA** and change the **Password**.

Enter the following CLI command to change the password to ha_pwd:

```
config system ha
    set password ha_pwd
end
```

Group ID

Similar to the group name, the group ID is also identifies the cluster. In most cases you do not have to change the group ID. However, you should change the group ID if you have more than one cluster on the same network. All members of the HA cluster must have the same group ID. The group ID is a number from 0 to 255.

Changing the group ID changes the cluster virtual MAC address. If two clusters on the same network have the same group ID you may encounter MAC address conflicts.

Enter the following CLI command to change the group ID to 10:

```
config system ha
    set group-id 10
end
```

Device failover, link failover, and session failover

The FGCP provides transparent device and link failover. You can also enable session pickup to provide session failover. A failover can be caused by a hardware failure, a software failure, or something as simple as a network cable being disconnected causing a link failover. When a failover occurs, the cluster detects and recognizes the failure and takes steps to respond so that the network can continue to operate without interruption. The internal operation of the cluster changes, but network components outside of the cluster notice little or no change.

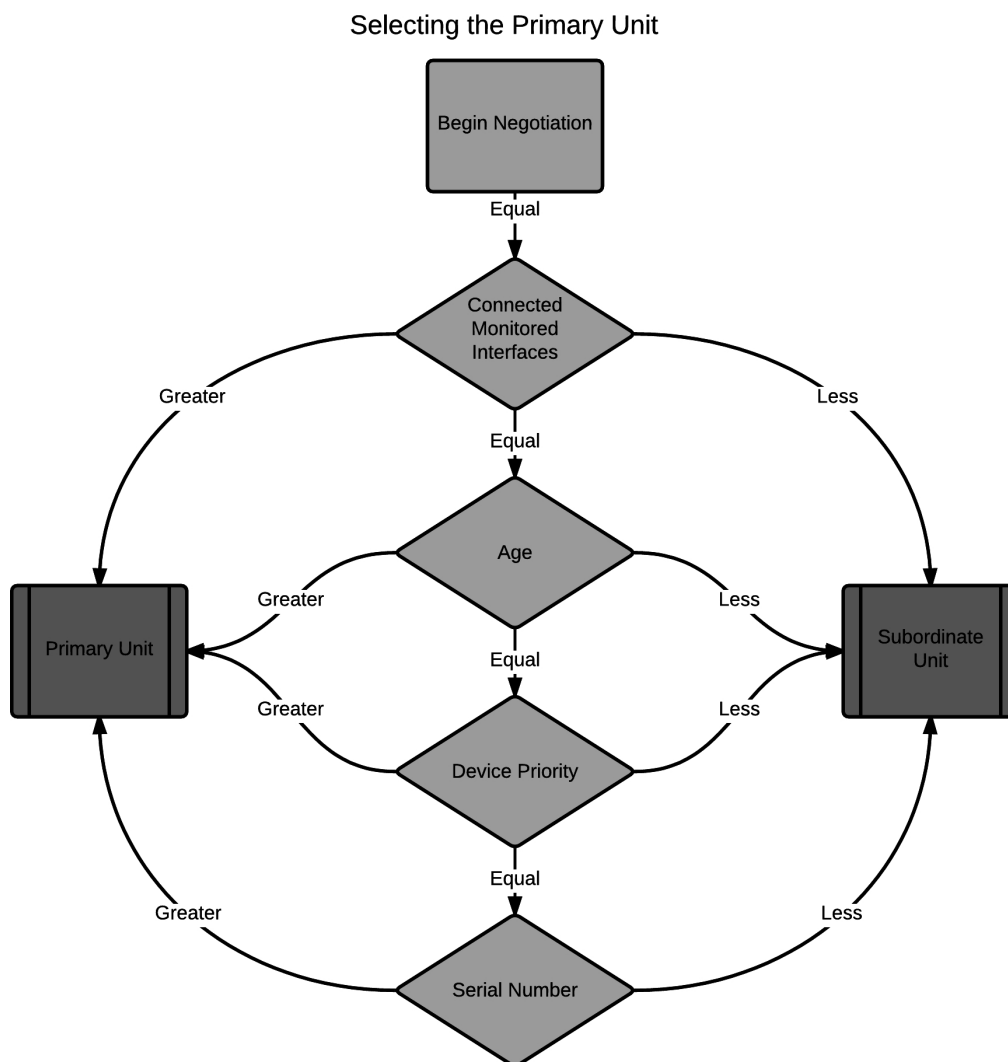
If a failover occurs, the cluster also records log messages about the event and can be configured to send log messages to a syslog server and to a FortiAnalyzer unit. The cluster can also send SNMP traps and alert email messages. These alerts can notify network administrators of the failover and may contain information that the network administrators can use to find and fix the problem that caused the failure.

For a complete description of device failover, link failover, and session failover, how clusters support these types of failover, and how FortiGate HA clusters compensate for a failure to maintain network traffic flow see [HA and failover protection on page 1521](#).

Primary unit selection

Once FortiGates recognize that they can form a cluster, the cluster units negotiate to select a primary unit. Primary unit selection occurs automatically based on the criteria shown below. After the cluster selects the primary unit, all of the remaining cluster units become subordinate units.

Negotiation and primary unit selection also takes place if a primary unit fails (device failover) or if a monitored interface fails or is disconnected (link failover). During a device or link failover, the cluster renegotiates to select a new primary unit also using the criteria shown below.



For many basic HA configurations primary unit selection simply selects the cluster unit with the highest serial number to become the primary unit. A basic HA configuration involves setting the HA mode to active-passive or

active-active and configuring the cluster group name and password. Using this configuration, the cluster unit with the highest serial number becomes the primary unit because primary unit selection disregards connected monitored interfaces (because interface monitoring is not configured), the age of the cluster units would usually always be the same, and all units would have the same device priority.

Using the serial number is a convenient way to differentiate cluster units; so basing primary unit selection on the serial number is predictable and easy to understand and interpret. Also the cluster unit with the highest serial number would usually be the newest FortiGate with the most recent hardware version. In many cases you may not need active control over primary unit selection, so basic primary unit selection based on serial number is sufficient.

In some situations you may want have control over which cluster unit becomes the primary unit. You can control primary unit selection by setting the device priority of one cluster unit to be higher than the device priority of all other cluster units. If you change one or more device priorities, during negotiation, the cluster unit with the highest device priority becomes the primary unit. As shown above, the FGCP selects the primary unit based on device priority before serial number. For more information about how to use device priorities, see [Primary unit selection and device priority on page 1390](#).

The only other way that you can influence primary unit selection is by configuring interface monitoring (also called port monitoring). Using interface monitoring you can make sure that cluster units with failed or disconnected monitored interfaces cannot become the primary unit. See [Primary unit selection and monitored interfaces on page 1387](#).

Finally, the age of a cluster unit is determined by a number of operating factors. Normally the age of all cluster units is the same so normally age has no effect on primary unit selection. Age does affect primary unit selection after a monitored interface failure. For more information about age, see [Primary unit selection and age on page 1387](#).

Viewing how the primary unit was selected

You can use the `get system ha status` command to see how the primary unit was selected. The output of this command contains a section called `Master selected using` that shows a history of how the primary unit was selected. For example, when a cluster first forms this part of the command output could have one line showing that the primary unit is the cluster unit with the highest uptime.

```
get system ha status
.
.
.
Master selected using:
    <2016/10/12 11:13:23> FG-5KD3914800344 is selected as the master because it
has the largest value of uptime.
.
.
.
```

Over time more messages could be added as the cluster negotiates to choose a new primary unit on different occasions. The command output below shows the cluster negotiated four times over a few days.

```
get system ha status
.
.
.
Master selected using:
```

```
<2016/10/16 11:36:07> FG-5KD3914800344 is selected as the master because it
has the largest value of uptime.
<2016/10/15 11:24:11> FG-5KD3914800284 is selected as the master because it
has the largest value of override priority.
<2016/10/13 11:15:13> FG-5KD3914800344 is selected as the master because it
has the largest value of uptime.
<2016/10/11 11:13:23> FG-5KD3914800344 is selected as the master because it
has the largest value of uptime.
.
.
.
```

Primary unit selection and monitored interfaces

If you have configured interface monitoring the cluster unit with the highest number of monitored interfaces that are connected to networks becomes the primary unit. Put another way, the cluster unit with the highest number of failed or disconnected monitored interfaces cannot become the primary unit.

Normally, when a cluster starts up, all monitored interfaces of all cluster units are connected and functioning normally. So monitored interfaces do not usually affect primary unit selection when the cluster first starts.

A cluster always renegotiates when a monitored interface fails or is disconnected (called link failover). A cluster also always renegotiates when a failed or disconnected monitored interface is restored.

If a primary unit monitored interface fails or is disconnected, the cluster renegotiates and if this is the only failed or disconnected monitored interface the cluster selects a new primary unit.

If a subordinate unit monitored interface fails or is disconnected, the cluster also renegotiates but will not necessarily select a new primary unit. However, the subordinate unit with the failed or disconnected monitored interface cannot become the primary unit.

Multiple monitored interfaces can fail or become disconnected on more than one cluster unit. Each time a monitored interface is disconnected or fails, the cluster negotiates to select the cluster unit with the most connected and operating monitored interfaces to become the primary unit. In fact, the intent of the link failover feature is just this, to make sure that the primary unit is always the cluster unit with the most connected and operating monitored interfaces.

Primary unit selection and age

The cluster unit with the highest age value becomes the primary unit. The age of a cluster unit is the amount of time since a monitored interface failed or is disconnected. Age is also reset when a cluster unit starts (boots up). So, when all cluster units start up at about the same time, they all have the same age. Age does not affect primary unit selection when all cluster units start up at the same time. Age also takes precedence over priority for primary unit selection.

If a link failure of a monitored interface occurs, the age value for the cluster unit that experiences the link failure is reset. So, the cluster unit that experienced the link failure also has a lower age value than the other cluster units. This reduced age does not effect primary unit selection because the number of link failures takes precedence over the age.

If the failed monitored interface is restored the cluster unit that had the failed monitored interface cannot become the primary unit because its age is still lower than the age of the other cluster units.

In most cases, the way that age is handled by the cluster reduces the number of times the cluster selects a new primary unit, which results in a more stable cluster since selecting a new primary unit has the potential to disrupt traffic.

Cluster age difference margin (grace period)

In any cluster, some of the cluster units may take longer to start up than others. This startup time difference can happen as a result of a number of issues and does not affect the normal operation of the cluster. To make sure that cluster units that start slower can still become primary units, by default the FGCP ignores age differences of up to 5 minutes (300 seconds).

In most cases, during normal operation this age difference margin or grace period helps clusters function as expected. However, the age difference margin can result in some unexpected behavior in some cases:

- During a cluster firmware upgrade with `uninterruptible-upgrade` enabled (the default configuration) the cluster should not select a new primary unit after the firmware of all cluster units has been updated. But since the age difference of the cluster units is most likely less than 300 seconds, age is not used to affect primary unit selection and the cluster may select a new primary unit.
- During failover testing where cluster units are failed over repeatedly the age difference between the cluster units will most likely be less than 5 minutes. During normal operation, if a failover occurs, when the failed unit rejoins the cluster its age will be very different from the age of the still operating cluster units so the cluster will not select a new primary unit. However, if a unit fails and is restored in a very short time the age difference may be less than 5 minutes. As a result the cluster may select a new primary unit during some failover testing scenarios.

Changing the cluster age difference margin

You can change the cluster age difference margin using the following command:

```
config system ha
    set ha-uptime-diff-margin 60
end
```

This command sets the cluster age difference margin to 60 seconds (1 minute). The age difference margin range is 1 to 65535 seconds. The default is 300 seconds.

You may want to reduce the margin if during failover testing you don't want to wait the default age difference margin of 5 minutes. You may also want to reduce the margin to allow uninterruptible upgrades to work. See [Operating clusters and virtual clusters on page 1487](#).

You may want to increase the age margin if cluster unit startup time differences are larger than 5 minutes.

Displaying cluster unit age differences

You can use the CLI command `diagnose sys ha dump-by group` to display the age difference of the units in a cluster. This command also displays information about a number of HA-related parameters for each cluster unit.

For example, consider a cluster of two FortiGate units. Entering the `diagnose sys ha dump-by group` command from the primary unit CLI displays information similar to the following:

```
diagnose sys ha dump-by group
    HA information.
group-id=0, group-name='External-HA-Cluster'

gmember_nr=2
'FGT6HD3916800525': ha_ip_idx=1, hb_packet_version=6, last_hb_jiffies=52097155,
```

```

linkfails=11, weight/o=0/0
    hbdev_nr=2: port3(mac=906c..70, last_hb_jiffies=52097155, hb_lost=0), port4(mac-
c=906c..71, last_hb_jiffies=52097155, hb_lost=0),
'FGT6HD3916801195': ha_ip_idx=0, hb_packet_version=6, last_hb_jiffies=0, link-
fails=0, weight/o=0/0

vcluster_nr=1
vcluster_0: start_time=1507754642(2017-10-11 13:44:02), state/o/chg_time=2(work)/2
(work)/1507754644(2017-10-11 13:44:04)
    'FGT6HD3916801955': ha_prio/o=1/1, link_failure=0(old=0), pingsvr_failure=0, flag-
g=0x00000000, uptime/reset_cnt=0/1
    'FGT6HD3916800525': ha_prio/o=0/0, link_failure=0(old=0), pingsvr_failure=0, flag-
g=0x00000001, uptime/reset_cnt=189/0

```

The last two lines of the output display status information about each cluster unit including the `uptime`. The `uptime` is the age difference in seconds/10 between the two units in the cluster.

In the example, the age of the subordinate unit 189 or is 18.9 seconds more than the age of the primary unit. The age difference is less than 5 minutes (less than 300 seconds) so age has no affect on primary unit selection. The cluster selected the unit with the highest serial number to be the primary unit.

If port1 (the monitored interface) of the primary unit is disconnected, the cluster renegotiates and the former subordinate unit becomes the primary unit. When you log into the new primary unit CLI and enter `diagnose sys ha dump-by group` you could get results similar to the following:

```

diagnose sys ha dump-by group
    HA information.
group-id=0, group-name='External-HA-Cluster'

gmember_nr=2
'FGT6HD3916800525': ha_ip_idx=1, hb_packet_version=6, last_hb_jiffies=52097155,
linkfails=11, weight/o=0/0
    hbdev_nr=2: port3(mac=906c..70, last_hb_jiffies=52097155, hb_lost=0), port4(mac-
c=906c..71, last_hb_jiffies=52097155, hb_lost=0),
'FGT6HD3916801195': ha_ip_idx=0, hb_packet_version=6, last_hb_jiffies=0, link-
fails=0, weight/o=0/0

vcluster_nr=1
vcluster_0: start_time=1507754642(2017-10-11 13:44:02), state/o/chg_time=2(work)/2
(work)/1507754644(2017-10-11 13:44:04)
    'FGT6HD3916800525': ha_prio/o=1/1, link_failure=0(old=0), pingsvr_failure=0, flag-
g=0x00000000, uptime/reset_cnt=0/1
    'FGT6HD3916801955': ha_prio/o=0/0, link_failure=0(old=0), pingsvr_failure=0, flag-
g=0x00000001, uptime/reset_cnt=1362/0

```

The command results show that the age of the new primary unit is 136.2 seconds higher than the age of the new subordinate unit.

If port1 of the former primary unit is reconnected the cluster will once again make this the primary unit because the age difference will still be less than 300 seconds. When you log into the primary unit CLI and enter `diagnose sys ha dump-by group` you get results similar to the following:

```

diagnose sys ha dump-by group
    HA information.
group-id=0, group-name='External-HA-Cluster'

```

```

gmember_nr=2
'FGT6HD3916800525': ha_ip_idx=1, hb_packet_version=6, last_hb_jiffies=52097155,
linkfails=11, weight/o=0/0
    hbdev_nr=2: port3(mac=906c..70, last_hb_jiffies=52097155, hb_lost=0), port4(mac-
c=906c..71, last_hb_jiffies=52097155, hb_lost=0),
'FGT6HD3916801195': ha_ip_idx=0, hb_packet_version=6, last_hb_jiffies=0, link-
fails=0, weight/o=0/0

vcluster_nr=1
vcluster_0: start_time=1507754642(2017-10-11 13:44:02), state/o/chg_time=2(work)/2
(work)/1507754644(2017-10-11 13:44:04)
    'FGT6HD3916800525': ha_prio/o=1/1, link_failure=0(old=0), pingsvr_failure=0, flag-
g=0x00000000, uptime/reset_cnt=0/1
    'FGT6HD3916801955': ha_prio/o=0/0, link_failure=0(old=0), pingsvr_failure=0, flag-
g=0x00000001, uptime/reset_cnt=-1362/0

```

Resetting the age of all cluster units

In some cases, age differences among cluster units can result in the wrong cluster unit or the wrong virtual cluster becoming the primary unit. For example, if a cluster unit set to a high priority reboots, that unit will have a lower age than other cluster units when it rejoins the cluster. Since age takes precedence over priority, the priority of this cluster unit will not be a factor in primary unit selection.

This problem also affects virtual cluster VDOM partitioning in a similar way. After a reboot of one of the units in a virtual cluster configuration, traffic for all VDOMs could continue to be processed by the cluster unit that did not reboot. This can happen because the age of both virtual clusters on the unit that did not reboot is greater than the age of both virtual clusters on the unit that rebooted.

One way to resolve this issue is to reboot all of the cluster units at the same time so that the age of all of the cluster units is reset. However, rebooting cluster units may interrupt or at least slow down traffic. If you would rather not reboot all of the cluster units you can instead use the following command to reset the age of individual cluster units.

```
diagnose sys ha reset-uptime
```

This command resets the age of a unit back to zero so that if no other unit in the cluster was reset at the same time, it will now have the lowest age. You would use this command to reset the age of the cluster unit that is currently the primary unit. Since it will have the lowest age, the other unit in the cluster will have the highest age and can then become the primary unit.



The `diagnose sys ha reset-uptime` command should only be used as a temporary solution. The command resets the HA age internally and does not affect the up time displayed for cluster units using the `diagnose sys ha dump-by all-vcluster` command or the up time displayed on the Dashboard or cluster members list. To make sure the actual up time for cluster units is the same as the HA age you should reboot the cluster units during a maintenance window.

Primary unit selection and device priority

A cluster unit with the highest device priority becomes the primary unit when the cluster starts up or renegotiates. By default, the device priority for all cluster units is 128. You can change the device priority to control which FortiGate becomes the primary unit during cluster negotiation. All other factors that influence primary unit selection either cannot be configured (age and serial number) or are synchronized among all cluster units.

(interface monitoring). You can set a different device priority for each cluster unit. During negotiation, if all monitored interfaces are connected, and all cluster units enter the cluster at the same time (or have the same age), the cluster with the highest device priority becomes the primary unit.

A higher device priority does not affect primary unit selection for a cluster unit with the most failed monitored interfaces or with an age that is higher than all other cluster units because failed monitored interfaces and age are used to select a primary unit before device priority.

Increasing the device priority of a cluster unit does not always guarantee that this cluster unit will become the primary unit. During cluster operation, an event that may affect primary unit selection may not always result in the cluster renegotiating. For example, when a unit joins a functioning cluster, the cluster will not renegotiate. So if a unit with a higher device priority joins a cluster the new unit becomes a subordinate unit until the cluster renegotiates.



Enabling the `override` HA CLI keyword makes changes in device priority more effective by causing the cluster to negotiate more often to make sure that the primary unit is always the unit with the highest device priority. For more information about `override`, see [Primary unit selection on page 1385](#).

Controlling primary unit selection by changing the device priority

You set a different device priority for each cluster unit to control the order in which cluster units become the primary unit when the primary unit fails.

To change the device priority from the GUI go to **System > HA** and change the **Device priority**.

Enter the following CLI command to change the device priority to 200:

```
config system ha
  set priority 200
end
```

The device priority is not synchronized among cluster units. In a functioning cluster you can change the device priority of any unit in the cluster. Whenever you change the device priority of a cluster unit, when the cluster negotiates, the unit with the highest device priority becomes the primary unit.

The following example shows how to change the device priority of a subordinate unit to 255 so that this subordinate unit becomes the primary unit. You can change the device priority of a subordinate unit by going to **System > HA** and selecting the Edit icon for the subordinate unit. Or from the CLI you can use the `execute ha manage 0` command to connect to the highest priority subordinate unit. After you enter the following commands the cluster renegotiates and selects a new primary unit.

```
execute ha manage 1
config system ha
  set priority 255
end
```

If you have three units in a cluster you can set the device priorities as shown below. When the cluster starts up, cluster unit A becomes the primary unit because it has the highest device priority. If unit A fails, unit B becomes the primary unit because unit B has a higher device priority than unit C.

Example device priorities for a cluster of three FortiGates

Cluster unit	Device priority
A	200
B	100
C	50

When configuring HA you do not have to change the device priority of any of the cluster units. If all cluster units have the same device priority, when the cluster first starts up the FGCP negotiates to select the cluster unit with the highest serial number to be the primary unit. Clusters also function normally if all units have the same device priority.

You can change the device priority if you want to control the roles that individual units play in the cluster. For example, if you want the same unit to always become the primary unit, set this unit device priority higher than the device priority of other cluster units. Also, if you want a cluster unit to always become a subordinate unit, set this cluster unit device priority lower than the device priority of other cluster units.

If you have a cluster of three units you can set a different priority for each unit to control which unit becomes the primary unit when all three cluster units are functioning and which will be the primary unit when two cluster units are functioning.

The device priority range is 0 to 255. The default device priority is 128.

If you are configuring a virtual cluster, if you have added virtual domains to both virtual clusters, you can set the device priority that the cluster unit has in virtual cluster 1 and virtual cluster 2. If a FortiGate has different device priorities in virtual cluster 1 and virtual cluster 2, the FortiGate may be the primary unit in one virtual cluster and the subordinate unit in the other.

Primary unit selection and the FortiGate serial number

The cluster unit with the highest serial number is more likely to become the primary unit. When first configuring FortiGates to be added to a cluster, if you do not change the device priority of any cluster unit, then the cluster unit with the highest serial number always becomes the primary unit.

Age does take precedence over serial number, so if a cluster unit takes longer to join a cluster for some reason (for example if one cluster unit is powered on after the others), that cluster unit will not become the primary unit because the other units have been in the cluster longer.

Device priority and failed monitored interfaces also take precedence over serial number. A higher device priority means a higher priority. So if you set the device priority of one unit higher or if a monitored interface fails, the cluster will not use the FortiGate serial number to select the primary unit.

Points to remember about primary unit selection

Some points to remember about primary unit selection:

- The FGCP compares primary unit selection criteria in the following order: Failed Monitored interfaces > Age > Device Priority > Serial number. The selection process stops at the first criteria that selects one cluster unit.
- Negotiation and primary unit selection is triggered if a cluster unit fails or if a monitored interface fails.

- If the HA age difference is more than 5 minutes (300 seconds), the cluster unit that is operating longer becomes the primary unit.
- If HA age difference is less than 5 minutes (300 seconds), the device priority and FortiGate serial number selects the cluster unit to become the primary unit.
- Every time a monitored interface fails the HA age of the cluster unit is reset to 0.
- Every time a cluster unit restarts the HA age of the cluster unit is reset to 0.

Temporarily setting a cluster unit to be the primary unit

You can use the following diagnose command to set a cluster unit to be the primary unit.

```
diagnose sys ha set-as-master enable
```



This command is intended for demonstration purposes and not for production use.
This command may not be visible for all FortiOS versions.

When you enter this command, the cluster immediately re-negotiates and the cluster unit on which you entered this command becomes the primary unit. This change is temporary and will be reverted if the cluster unit restarts.

You can also use the following command from the same cluster unit to turn this option off, causing the cluster to renegotiate and select a new primary unit.

```
diagnose sys ha set-as-master disable
```

You can also configure when to disabling the set-as-master setting. For example, to disable the set as master setting on January 25, 2015 you can enter a date after the disable keyword:

```
diagnose sys ha set-as-master disable 2015 01 25
```

HA override

The HA `override` CLI keyword is disabled by default. When `override` is disabled a cluster may not always renegotiate when an event occurs that affects primary unit selection. For example, when `override` is disabled a cluster will not renegotiate when you change a cluster unit device priority or when you add a new cluster unit to a cluster. This is true even if the unit added to the cluster has a higher device priority than any other unit in the cluster. Also, when `override` is disabled a cluster does not negotiate if the new unit added to the cluster has a failed or disconnected monitored interface.



For a virtual cluster configuration, `override` is enabled by default for both virtual clusters when you enable virtual cluster 2. For more information, see [Virtual clusters on page 1463](#).

In most cases you should keep `override` disabled to reduce how often the cluster negotiates. Frequent negotiations may cause frequent traffic interruptions.

However, if you want to make sure that the same cluster unit always operates as the primary unit and if you are less concerned about frequent cluster negotiation you can set its device priority higher than other cluster units and enable `override`.

To enable `override`, connect to each cluster unit CLI (using the `execute ha manage` command) and use the `config system ha` CLI command to enable `override`.

For `override` to be effective, you must also set the device priority highest on the cluster unit that you want to always be the primary unit. To increase the device priority, from the CLI use the `config system ha` command and increase the value of the `priority` keyword to a number higher than the default priority of 128.

You can also increase the device priority from the GUI by going to **System > HA**. To increase the device priority of the primary unit select edit for the primary or subordinate unit and set the **Device Priority** to a number higher than 128.



The `override` setting and device priority value are not synchronized to all cluster units. You must enable override and adjust device priority manually and separately for each cluster unit.

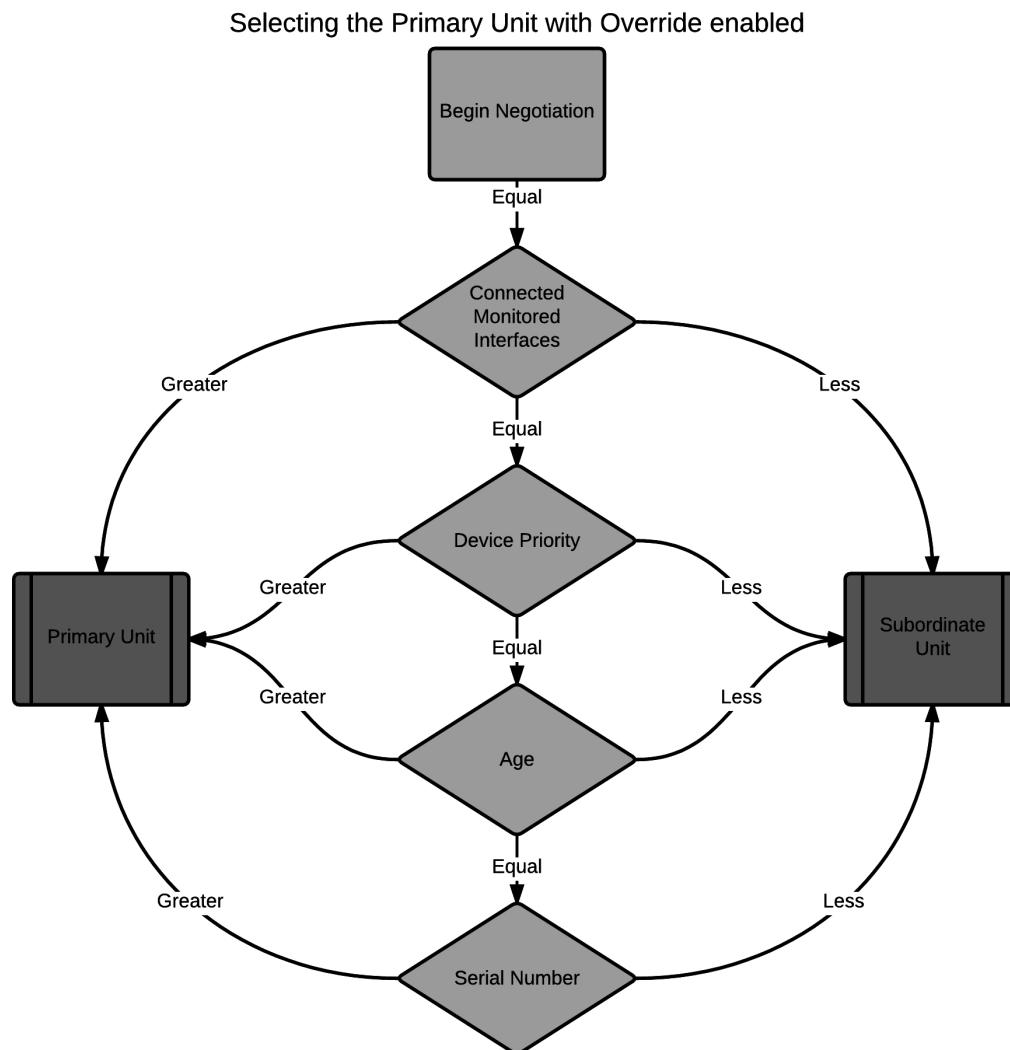
With `override` enabled, the primary unit with the highest device priority will always become the primary unit. Whenever an event occurs that may affect primary unit selection, the cluster negotiates. For example, when `override` is enabled a cluster renegotiates when you change the device priority of any cluster unit or when you add a new unit to a cluster.

Override and primary unit selection

Enabling `override` changes the order of primary unit selection. As shown below, if `override` is enabled, primary unit selection considers device priority before age and serial number. This means that if you set the device priority higher on one cluster unit, with `override` enabled this cluster unit becomes the primary unit even if its age and serial number are lower than other cluster units.

Similar to when `override` is disabled, when `override` is enabled primary unit selection checks for connected monitored interfaces first. So if interface monitoring is enabled, the cluster unit with the most disconnected monitored interfaces cannot become the primary unit, even if the unit has the highest device priority.

If all monitored interfaces are connected (or interface monitoring is not enabled) and the device priority of all cluster units is the same then age and serial number affect primary unit selection.



Controlling primary unit selection using device priority and override

To configure one cluster unit to always become the primary unit you should set its device priority to be higher than the device priorities of the other cluster units and you should enable `override` on all cluster units.

Using this configuration, when the cluster is operating normally the primary unit is always the unit with the highest device priority. If the primary unit fails the cluster renegotiates to select another cluster unit to be the primary unit. If the failed primary unit recovers, starts up again and rejoins the cluster, because `override` is enabled, the cluster renegotiates. Because the restarted primary unit has the highest device priority it once again becomes the primary unit.

In the same situation with `override` disabled, because the age of the failed primary unit is lower than the age of the other cluster units, when the failed primary unit rejoins the cluster it does not become the primary unit. Instead, even though the failed primary unit may have the highest device priority it becomes a subordinate unit because its age is lower than the age of all the other cluster units.

Points to remember about primary unit selection when override is enabled

Some points to remember about primary unit selection when `override` is enabled:

- The FGCP compares primary unit selection criteria in the following order: Failed Monitored Interfaces > Device Priority > Age > Serial number. The selection process stops at the first criteria that selects one cluster unit.
- Negotiation and primary unit selection is triggered whenever an event occurs which may affect primary unit selection. For example negotiation occurs, when you change the device priority, when you add a new unit to a cluster, if a cluster unit fails, or if a monitored interface fails.
- Device priority is considered before age. Otherwise age is handled the same when `override` is enabled.

Configuration changes can be lost if override is enabled

In some cases, when `override` is enabled and you make configuration changes to an HA cluster these changes can be lost. For example, consider the following sequence:

1. A cluster of two FortiGates is operating with `override` enabled.
 - FGT-A: Primary unit with device priority 200 and with `override` enabled
 - FGT-B: Subordinate unit with device priority 100 and with `override` disabled
 - If both units are operating, FGT-A always becomes the primary unit because FGT-A has the highest device priority.
2. FGT-A fails and FGT-B becomes the new primary unit.
3. The administrator makes configuration changes to the cluster.

The configuration changes are made to FGT-B because FGT-B is operating as the primary unit. These configuration changes are not synchronized to FGT-A because FGT-A is not operating.
4. FGT-A is restored and starts up again.
5. The cluster renegotiates and FGT-A becomes the new primary unit.
6. The cluster recognizes that the configurations of FGT-A and FGT-B are not the same.
7. The configuration of FGT-A is synchronized to FGT-B.

The configuration is always synchronized from the primary unit to the subordinate units.
8. The cluster is now operating with the same configuration as FGT-A. The configuration changes made to FGT-B have been lost.

The solution

When `override` is enabled, you can prevent configuration changes from being lost by doing the following:

- Verify that all cluster units are operating before making configuration changes (from the GUI go to **System > HA** to view the cluster members list or from the FortiOS CLI enter `get system ha status`).
- Make sure the device priority of the primary unit is set higher than the device priorities of all other cluster units before making configuration changes.
- Disable `override` either permanently or until all configuration changes have been made and synchronized to all cluster units.

Override and disconnecting a unit from a cluster

A similar scenario to that described above may occur when `override` is enabled and you use the Disconnect from Cluster option from the GUI or the `execute ha disconnect` command from the CLI to disconnect a cluster unit from a cluster.

Configuration changes made to the cluster can be lost when you reconnect the disconnected unit to the cluster. You should make sure that the device priority of the disconnected unit is lower than the device priority of the current primary unit. Otherwise, when the disconnected unit joins the cluster, if `override` is enabled, the cluster renegotiates and the disconnected unit may become the primary unit. If this happens, the configuration of the disconnected unit is synchronized to all other cluster units and any configuration changes made between when the unit was disconnected and reconnected are lost.

Delaying how quickly the primary unit rejoins the cluster when override is enabled

In some cases when `override` is enabled and the unit designated to be the primary unit rejoins the cluster it will become the primary unit too soon and cause traffic disruption. This can happen, for example, if one of the FortiGate interfaces gets its address using PPPoE. If the backup unit is operating as the primary unit and processing traffic, when the primary unit comes up it may need a short time to get a new IP address from the PPPoE server. If the primary unit takes over the cluster before it has an IP address, traffic will be disrupted until the primary unit gets its address.

You can resolve this problem by using the following command to add a wait time. In this example the wait time is 10 seconds. The wait time range is 0 to 3600 seconds and the default wait time is 0 seconds.

```
config system ha
    set override-wait-time 10
end
```

With this wait time configured, after the primary unit is up and running it has 10 seconds to synchronize sessions, get IP address(es) from PPPoE and DHCP servers and so on. After 10 seconds the primary unit sends gratuitous arp packets and all traffic to the cluster is sent to the new primary unit. You can adjust the wait time according to the conditions on your network.

FortiGate HA compatibility with DHCP and PPPoE

FortiGate HA is compatible with DHCP and PPPoE but care should be taken when configuring a cluster that includes a FortiGate interface configured to get its IP address with DHCP or PPPoE. Fortinet recommends that you turn on DHCP or PPPoE addressing for an interface after the cluster has been configured. If an interface is configured for DHCP or PPPoE, turning on high availability may result in the interface receiving an incorrect address or not being able to connect to the DHCP or PPPoE server correctly.



You cannot switch to operate in HA mode if one or more FortiGate interfaces is configured as a PPTP or L2TP client.

You can configure a cluster to act as a DHCP server or a DHCP relay agent. In both active-passive and active-active clusters DHCP relay sessions are always handled by the primary unit. It is possible that a DHCP relay session could be interrupted by a failover. If this occurs the DHCP relay session is not resumed after the failover and the DHCP client may have to repeat the DHCP request.

When a cluster is operating as a DHCP server the primary unit responds to all DHCP requests and maintains the DHCP server address lease database. The cluster also dynamically synchronizes the DHCP server address lease database to the subordinate units. If a failover occurs, the new primary unit will have an up-to-date DHCP server address lease database. Synchronizing the DHCP address lease database prevents the new primary unit from responding incorrectly to new DHCP requests after a failover.

Also, it is possible that when FortiGates first negotiate to form a cluster that a unit that ends up as a subordinate unit in the cluster will have information in its DHCP address lease database that the cluster unit operating as the

primary unit does not have. This can happen if a FortiGate responds to DHCP requests while operating as a standalone unit and then when the cluster is formed this unit becomes a subordinate unit. Because of this possibility, after a cluster is formed the DHCP address lease databases of all of the cluster units are merged into one database which is then synchronized to all cluster units.


HA and distributed clustering

The FGCP supports widely separated cluster units installed in different physical locations. Distributed clusters can have cluster units in different rooms in the same building, different buildings in the same location, or even different geographical sites such as different cities, countries or continents.

Just like any cluster, distributed clusters require heartbeat communication between cluster units. In a distributed cluster this heartbeat communication can take place over the internet or over other transmission methods including satellite linkups.

Most Data Center Interconnect (DCI) or MPLS-based solutions that support layer 2 extensions between the remote data centers should also support HA heartbeat communication between the FortiGates in the distributed locations. Using VLANs and switches in promiscuous mode to pass all traffic between the locations can also be helpful.

HA heartbeat IP addresses are not configurable so the heartbeat interfaces have to be able to communicate over the same subnet. See [HA heartbeat interface IP addresses on page 1526](#).



Because of the possible distance it may take a relatively long time for heartbeat packets to be transmitted between cluster units. This could lead to a split brain scenario. To avoid a split brain scenario you can increase the heartbeat interval so that the cluster expects extra time between heartbeat packets. A general rule is to configure the failover time to be longer than the max latency. You could also increase the `hb-lost-threshold` to tolerate losing heartbeat packets if the network connection is less reliable.

In addition you could use different link paths for heartbeat packets to optimize HA heartbeat communication. You could also configure QoS on the links used for HA heartbeat traffic to make sure heartbeat communication has the highest priority.

For information about changing the heartbeat interval and other heartbeat related settings, see [Modifying heartbeat timing on page 1528](#).

Clusters of three or four FortiGates

The FGCP supports a cluster of two, three, or four FortiGates. You can add more than two units to a cluster to improve reliability: if two cluster units fail the third will continue to operate and so on. A cluster of three or four units in active-active mode may improve performance since another cluster unit is available for security profile processing. However, active-active FGCP HA results in diminishing performance returns as you add units to the cluster, so the additional performance achieved by adding the third cluster unit may not be worth the cost.

There are no special requirements for clusters of more than two units. Here are a few recommendations though:

- The matching heartbeat interfaces of all of the cluster units must be able to communicate with each other. So each unit's matching heartbeat interface should be connected to the same switch. If the ha1 interface is used for heartbeat communication, then the ha1 interfaces of all of the units in the cluster must be connected together so communication can happen between all of the cluster units over the ha1 interface.
- Redundant heartbeat interfaces are recommended. You can reduce the number of points of failure by connecting each matching set of heartbeat interfaces to a different switch. This is not a requirement; however, and you can connect both heartbeat interfaces of all cluster units to the same switch. However, if that switch fails the cluster will stop forwarding traffic.
- For any cluster, a dedicated switch for each heartbeat interface is recommended because of the large volume of heartbeat traffic and to keep heartbeat traffic off of other networks, but it is not required.
- Full mesh HA can scale to three or four FortiGates. Full mesh HA is not required if you have more than 2 units in a cluster.
- Virtual clustering can only be done with two FortiGates.

Connecting a cluster of three FortiGates

This example shows how to connect a cluster of three FortiGates where:

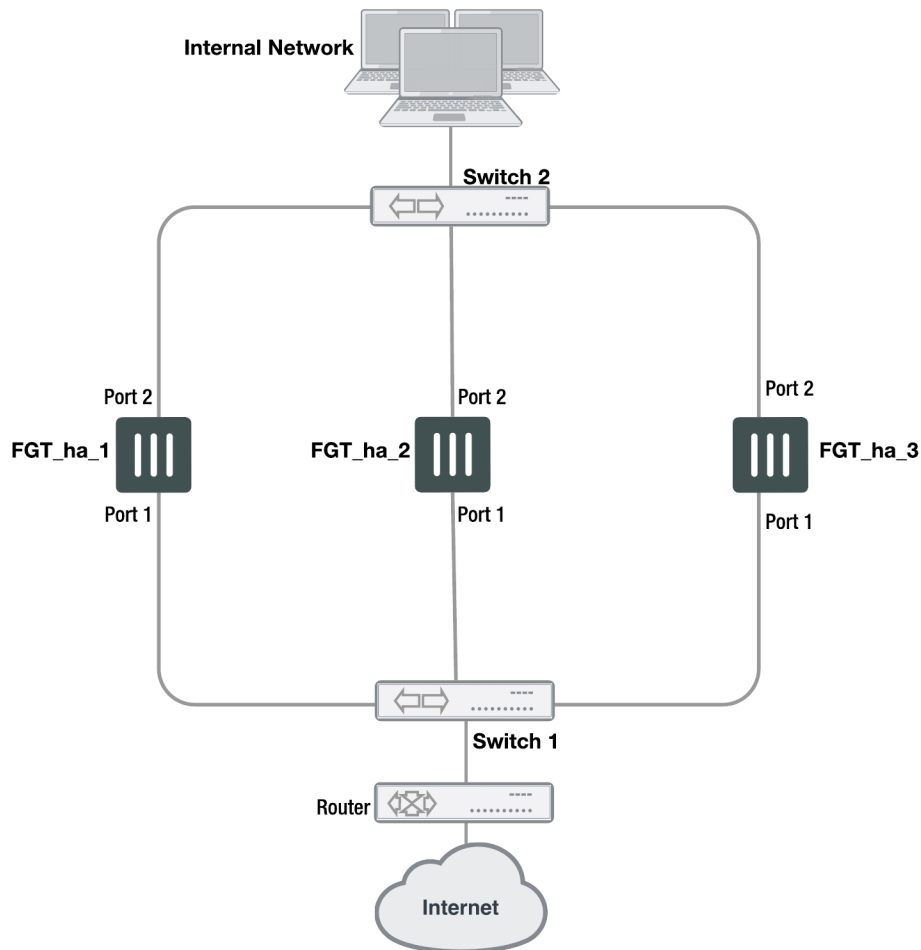
- Port1 connects the cluster to the internet
- Port2 connects the cluster to the internal network
- Port3 and Port4 are the heartbeat interfaces

Use the following steps to connect the cluster units to each other and to their networks:

1. Connect the network interfaces:

- Connect the port1 interface of each FortiGate to the same switch (Switch 1) and connect this switch to the internet.
- Connect the port2 interface of each FortiGate to the same switch (Switch 2) and connect this switch to the internal Network.


Connecting the network interfaces (cluster of three FortiGates)



2. Connect the heartbeat interfaces:

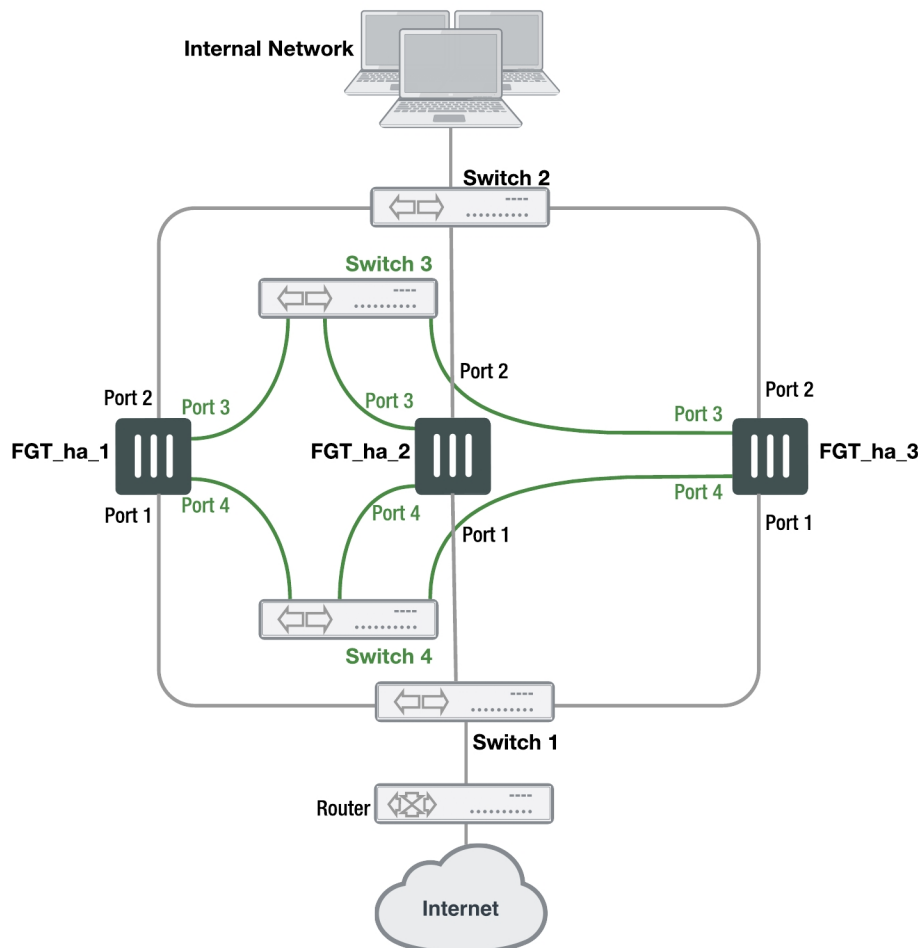
- Connect the port3 interface of each FortiGate to the same switch (Switch 3)
- Connect the port4 interface of each FortiGate to the same switch (Switch 4)

Connecting the heartbeat interfaces (cluster of three FortiGates)



The network and heartbeat connections when combined into one diagram appear like the following:

Network and heartbeat interface connections (cluster of three FortiGates)



Disk storage configuration and HA

If your cluster units include storage disks (for example for storing log messages, WAN optimization data and web caching) all cluster units must have identical storage disk configurations. This means each cluster unit must have same number of disks (including AMC and FortiGate Storage Module (FSM) hard disks) and also means that matching disks in each cluster unit must be the same size, have the same format, and have the same number of partitions.

In most cases the default hard disk configuration of the cluster units will be compatible. However, a hard disk formatted by an older FortiGate firmware version may not be compatible with a hard disk formatted by a more recent firmware version. Problems may also arise if you have used the `execute scsi-dev` command to add or change hard disk protections.

If a cluster unit CLI displays hard disk compatibility messages, you may need to use the `execute scsi-dev delete` command to delete partitions. You can also use the `execute formatlogdisk` command to

reformat disks. In some cases after deleting all partitions and reformatting the disks, you may still see disk incompatibility messages. If this happens, visit the [Fortinet Support](#) website assistance.

FGCP high availability best practices

Fortinet suggests the following practices related to high availability:

- Use Active-Active HA to distribute TCP and UTM sessions among multiple cluster units. An active-active cluster may have higher throughput than a standalone FortiGate unit or than an active-passive cluster.
- Use a different host name on each FortiGate unit when configuring an HA cluster. Fewer steps are required to add host names to each cluster unit before configuring HA and forming a cluster.
- Consider adding an Alias to the interfaces used for the HA heartbeat so that you always get a reminder about what these interfaces are being used for.
- Enabling `load-balance-all` can increase device and network load since more traffic is load-balanced. This may be appropriate for use in a deployment using the firewall capabilities of the FortiGate unit and IPS but no other content inspection.
- An advantage of using session pickup is that non-content inspection sessions will be picked up by the new primary unit after a failover. The disadvantage is that the cluster generates more heartbeat traffic to support session pickup as a larger portion of the session table must be synchronized. Session pickup should be configured only when required and is not recommended for use with SOHO FortiGate models. Session pickup should only be used if the primary heartbeat link is dedicated (otherwise the additional HA heartbeat traffic could affect network performance).
- If session pickup is not selected, after a device or link failover all sessions are briefly interrupted and must be re-established at the application level after the cluster renegotiates. For example, after a failover, users browsing the web can just refresh their browsers to resume browsing. Users downloading large files may have to restart their download after a failover. Other protocols may experience data loss and some protocols may require sessions to be manually restarted. For example, a user downloading files with FTP may have to either restart downloads or restart their FTP client.
- If you need to enable session pickup, consider enabling `session-pickup-delay` to improve performance by reducing the number of sessions that are synchronized. See [Improving session synchronization performance on page 1](#).
- Consider using the `session-sync-dev` option to move session synchronization traffic off the HA heartbeat link to one or more dedicated session synchronization interfaces. See [Improving session synchronization performance on page 1](#).
- To avoid unpredictable results, when you connect a switch to multiple redundant or aggregate interfaces in an active-passive cluster you should configure separate redundant or aggregate interfaces on the switch; one for each cluster unit.
- Use SNMP, syslog, or email alerts to monitor a cluster for failover messages. Alert messages about cluster failovers may help find and diagnose network problems quickly and efficiently.

Heartbeat interfaces

Fortinet suggests the following practices related to heartbeat interfaces:



Do not use a FortiGate switch port for the HA heartbeat traffic. This configuration is not supported.

- Configure at least two heartbeat interfaces and set these interfaces to have different priorities.
- For clusters of two FortiGate units, as much as possible, heartbeat interfaces should be directly connected using patch cables (without involving other network equipment such as switches). If switches have to be used they should not be used for other network traffic that could flood the switches and cause heartbeat delays.
 - If you cannot use a dedicated switch, the use of a dedicated VLAN can help limit the broadcast domain to protect the heartbeat traffic and the bandwidth it creates.
- For clusters of three or four FortiGate units, use switches to connect heartbeat interfaces. The corresponding heartbeat interface of each FortiGate unit in the cluster must be connected to the same switch. For improved redundancy use a different switch for each heartbeat interface. In that way if the switch connecting one of the heartbeat interfaces fails or is unplugged, heartbeat traffic can continue on the other heartbeat interfaces and switch.
- Isolate heartbeat interfaces from user networks. Heartbeat packets contain sensitive cluster configuration information and can consume a considerable amount of network bandwidth. If the cluster consists of two FortiGate units, connect the heartbeat interfaces directly using a crossover cable or a regular Ethernet cable. For clusters with more than two units, connect heartbeat interfaces to a separate switch that is not connected to any network.
- If heartbeat traffic cannot be isolated from user networks, enable heartbeat message encryption and authentication to protect cluster information. See [Enabling or disabling HA heartbeat encryption and authentication on page 1529](#).
- Configure and connect redundant heartbeat interfaces so that if one heartbeat interface fails or becomes disconnected, HA heartbeat traffic can continue to be transmitted using the backup heartbeat interface. If heartbeat communication fails, all cluster members will think they are the primary unit resulting in multiple devices on the network with the same IP addresses and MAC addresses (condition referred to as *Split Brain*) and communication will be disrupted until heartbeat communication can be reestablished.
- Do not monitor dedicated heartbeat interfaces; monitor those interfaces whose failure should trigger a device failover.
- Where possible at least one heartbeat interface should not be connected to an NP4 or NP6 processor to avoid NP4 or NP6-related problems from affecting heartbeat traffic.
- Where possible, the heartbeat interfaces should not be connected to an NP4 or NP6 processor that is also processing network traffic.
- Where possible, each heartbeat interface should be connected to a different NP4 or NP6 processor.
- Any FortiGate interface can be used as a heartbeat interface including 10/100/1000Base-T, SFP, QSFP fiber and copper, and so on. If you set up two or more interfaces as heartbeat interfaces each interface can be a different type and speed.

Interface monitoring (port monitoring)

Fortinet suggests the following practices related to interface monitoring (also called port monitoring):

- Wait until a cluster is up and running and all interfaces are connected before enabling interface monitoring. A monitored interface can easily become disconnected during initial setup and cause failovers to occur before the cluster is fully configured and tested.
- Monitor interfaces connected to networks that process high priority traffic so that the cluster maintains connections to these networks if a failure occurs.
- Avoid configuring interface monitoring for all interfaces.
- Supplement interface monitoring with remote link failover. Configure remote link failover to maintain packet flow if a link not directly connected to a cluster unit (for example, between a switch connected to a cluster interface and the network) fails. See [Remote link failover on page 1556](#).

FGCP HA terminology

The following HA-specific terms are used in this document.

Cluster

A group of FortiGates that act as a single virtual FortiGate to maintain connectivity even if one of the FortiGates in the cluster fails.

Cluster unit

A FortiGate operating in a FortiGate HA cluster.

Device failover

Device failover is a basic requirement of any highly available system. Device failover means that if a device fails, a replacement device automatically takes the place of the failed device and continues operating in the same manner as the failed device.

Failover

A FortiGate taking over processing network traffic in place of another unit in the cluster that suffered a device failure or a link failure.

Failure

A hardware or software problem that causes a FortiGate or a monitored interface to stop processing network traffic.

FGCP

The FortiGate clustering protocol (FGCP) that specifies how the FortiGates in a cluster communicate to keep the cluster operating.

Full mesh HA

Full mesh HA is a method of removing single points of failure on a network that includes an HA cluster. FortiGate models that support redundant interfaces can be used to create a cluster configuration called full mesh HA. Full mesh HA includes redundant connections between all network components. If any single component or any single connection fails, traffic switches to the redundant component or connection.

HA virtual MAC address

When operating in HA mode, all of the interfaces of the primary unit acquire the same HA virtual MAC address. All communications with the cluster must use this MAC address. The HA virtual MAC address is set according to the group ID.

Heartbeat

Also called FGCP heartbeat or HA heartbeat. The heartbeat constantly communicates HA status and synchronization information to make sure that the cluster is operating properly.

Heartbeat device

An Ethernet network interface in a cluster that is used by the FGCP for heartbeat communications among cluster units.

Heartbeat failover

If an interface functioning as the heartbeat device fails, the heartbeat is transferred to another interface also configured as an HA heartbeat device.

Hello state

In the hello state a cluster unit has powered on in HA mode, is using HA heartbeat interfaces to send hello packets, and is listening on its heartbeat interfaces for hello packets from other FortiGates. Hello state may appear in HA log messages.

High availability

The ability that a cluster has to maintain a connection when there is a device or link failure by having another unit in the cluster take over the connection, without any loss of connectivity. To achieve high availability, all FortiGates in the cluster share session and configuration information.

Interface monitoring

You can configure interface monitoring (also called port monitoring) to monitor FortiGate interfaces to verify that the monitored interfaces are functioning properly and connected to their networks. If a monitored interface fails or is disconnected from its network the interface leaves the cluster and a link failover occurs. For more information about interface monitoring, see [Link failover \(port monitoring or interface monitoring\)](#) on page 1549.

Link failover

Link failover means that if a monitored interface fails, the cluster reorganizes to re-establish a link to the network that the monitored interface was connected to and to continue operating with minimal or no disruption of network traffic.

Load balancing

Also known as active-active HA. All units in the cluster process network traffic. The FGCP employs a technique similar to unicast load balancing. The primary unit interfaces are assigned virtual MAC addresses which are associated on the network with the cluster IP addresses. The primary unit is the only cluster unit to receive packets sent to the cluster. The primary unit can process packets itself, or propagate them to subordinate units according to a load balancing schedule. Communication between the cluster units uses the actual cluster unit MAC addresses.

Monitored interface

An interface that is monitored by a cluster to make sure that it is connected and operating correctly. The cluster monitors the connectivity of this interface for all cluster units. If a monitored interface fails or becomes disconnected from its network, the cluster will compensate.

Primary unit

Also called the primary cluster unit, this cluster unit controls how the cluster operates. The primary unit sends hello packets to all cluster units to synchronize session information, synchronize the cluster configuration, and to synchronize the cluster routing table. The hello packets also confirm for the subordinate units that the primary unit is still functioning.

The primary unit also tracks the status of all subordinate units. When you start a management connection to a cluster, you connect to the primary unit.

In an active-passive cluster, the primary unit processes all network traffic. If a subordinate unit fails, the primary unit updates the cluster configuration database.

In an active-active cluster, the primary unit receives all network traffic and re-directs this traffic to subordinate units. If a subordinate unit fails, the primary unit updates the cluster status and redistributes load balanced traffic to other subordinate units in the cluster.

The FortiGate firmware uses the term master to refer to the primary unit.

Session failover

Session failover means that a cluster maintains active network sessions after a device or link failover. FortiGate HA does not support session failover by default. To enable session failover you must change the HA configuration to select Enable Session Pick-up.

Session pickup

If you enable session pickup for a cluster, if the primary unit fails or a subordinate unit in an active-active cluster fails, all communication sessions with the cluster are maintained or picked up by the cluster after the cluster negotiates to select a new primary unit.

If session pickup is not a requirement of your HA installation, you can disable this option to save processing resources and reduce the network bandwidth used by HA session synchronization. In many cases interrupted sessions will resume on their own after a failover even if session pickup is not enabled. You can also enable session pickup delay to reduce the number of sessions that are synchronized by session pickup.

Standby state

A subordinate unit in an active-passive HA cluster operates in the standby state. In a virtual cluster, a subordinate virtual domain also operates in the standby state. The standby state is actually a hot-standby state because the subordinate unit or subordinate virtual domain is not processing traffic but is monitoring the primary unit session table to take the place of the primary unit or primary virtual domain if a failure occurs.

In an active-active cluster all cluster units operate in a work state.

When standby state appears in HA log messages this usually means that a cluster unit has become a subordinate unit in an active-passive cluster or that a virtual domain has become a subordinate virtual domain.

State synchronization

The part of the FGCP that maintains connections after failover.

Subordinate unit

Also called the subordinate cluster unit, each cluster contains one or more cluster units that are not functioning as the primary unit. Subordinate units are always waiting to become the primary unit. If a subordinate unit does not receive hello packets from the primary unit, it attempts to become the primary unit.

In an active-active cluster, subordinate units keep track of cluster connections, keep their configurations and routing tables synchronized with the primary unit, and process network traffic assigned to them by the primary unit. In an active-passive cluster, subordinate units do not process network traffic. However, active-passive subordinate units do keep track of cluster connections and do keep their configurations and routing tables synchronized with the primary unit.

The FortiGate firmware uses the terms slave and subsidiary unit to refer to a subordinate unit.

Virtual clustering

Virtual clustering is an extension of the FGCP for FortiGates operating with multiple VDOMS enabled. Virtual clustering operates in active-passive mode to provide failover protection between two instances of a VDOM operating on two different cluster units. You can also operate virtual clustering in active-active mode to use HA load balancing to load balance sessions between cluster units. Alternatively, by distributing VDOM processing between the two cluster units you can also configure virtual clustering to provide load balancing by distributing sessions for different VDOMs to each cluster unit.

Work state

The primary unit in an active-passive HA cluster, a primary virtual domain in a virtual cluster, and all cluster units in an active-active cluster operate in the work state. A cluster unit operating in the work state processes traffic, monitors the status of the other cluster units, and tracks the session table of the cluster.

When work state appears in HA log messages this usually means that a cluster unit has become the primary unit or that a virtual domain has become a primary virtual domain.

FGCP support for OCVPN

You can set up a One-click VPN (OCVPN) on an FGCP cluster without any special configuration steps. When you add an OCVPN configuration, the FGCP synchronizes the configuration to all of the FortiGates in the cluster. When an OCVPN tunnel comes up between a remote client and the cluster, the OCVPN communicates with the primary FortiGate. The FGCP then synchronizes the VPN sessions to the other FortiGates in the cluster. If a failover occurs, the OCVPN sessions fail over to the new primary FortiGate and the OCVPN sessions continue with only minor interruptions.

A standalone FortiGate OCVPN configuration is not compatible with an FGCP OCVPN configuration. If you set up OCVPN on a stand-alone FortiGate, before you add this stand-alone FortiGate to an FGCP cluster you must disable any OCVPN configurations, set up HA, and then re-create the OCVPN configurations after the cluster is established.

The reverse is also true. If you decide to convert a cluster with an OCVPN configuration to a stand-alone FortiGate, you need to remove the OCVPN configuration, set up the standalone FortiGate and then re-create the OCVPN configuration on the standalone FortiGate.

HA GUI options

Go to **System > HA** to change HA options. You can set the following options to put a FortiGate into HA mode. You can also change any of these options while the cluster is operating.

You can configure HA options for a FortiGate with virtual domains (VDOMs) enabled by logging into the GUI as the global admin administrator and going to **System > HA**.

If already operating in HA mode, go to **System > HA** to display the cluster members list. You can then edit the primary unit to change HA settings.

Go to **System > HA > View HA Statistics** to view statistics about cluster operation.



Most virtual cluster HA options are the same as normal HA options. However, virtual clusters include VDOM partitioning options. Other differences between configuration options for regular HA and for virtual clustering HA are described below and see [Virtual clusters on page 1463](#).



FortiGate HA is compatible with DHCP and PPPoE but care should be taken when configuring a cluster that includes a FortiGate interface configured to get its IP address with DHCP or PPPoE. Fortinet recommends that you turn on DHCP or PPPoE addressing for an interface after the cluster has been configured. If an interface is configured for DHCP or PPPoE, turning on high availability may result in the interface receiving an incorrect address or not being able to connect to the DHCP or PPPoE server correctly.

Mode

Select an HA mode for the cluster or return the FortiGate in the cluster to standalone mode. When configuring a cluster, you must set all members of the HA cluster to the same HA mode. You can select **Standalone** (to disable HA), **Active-Passive**, or **Active-Active**.

If virtual domains are enabled you can select **Active-Passive** or **Standalone**.

Device priority

Optionally set the device priority of the cluster FortiGate. Each FortiGate in a cluster can have a different device priority. During HA negotiation, the FortiGate with the highest device priority usually becomes the primary unit.

In a virtual cluster configuration, each cluster FortiGate can have two different device priorities, one for each virtual cluster. During HA negotiation, the FortiGate with the highest device priority in a virtual cluster becomes the primary FortiGate for that virtual cluster.

Changes to the device priority are not synchronized. You can accept the default device priority when first configuring a cluster.

Synchronize management VDOM

This options appears if you have enabled multiple VDOMS and set a VDOM other than the root VDOM to be the management VDOM. You can disable this option to prevent the management VDOM configuration from being synchronized between cluster units in the virtual cluster. This allows you to add an interface to the VDOM in each

cluster unit and then to give the Interface a different IP address in each cluster unit, allowing you to manage each cluster unit separately.

You can also enable this feature using the following command:

```
config system ha
    set standalone-mgmt-vdom enable
end
```

Group name

Enter a name to identify the cluster. The maximum length of the group name is 32 characters. The group name must be the same for all cluster units before the cluster units can form a cluster. After a cluster is operating, you can change the group name. The group name change is synchronized to all cluster units.

Password

Enter a password to identify the cluster. The password must be the same for all cluster FortiGates before the cluster FortiGates can form a cluster.

Two clusters on the same network must have different passwords.

The password is synchronized to all cluster units in an operating cluster. If you change the password of one cluster unit the change is synchronized to all cluster units.

Session pickup

Select to enable session pickup so that if the primary unit fails, sessions are picked up by the cluster unit that becomes the new primary unit.

You must enable session pickup for session failover protection. If you do not require session failover protection, leaving session pickup disabled may reduce HA CPU usage and reduce HA heartbeat network bandwidth usage. See [Session failover \(session pick-up\) on page 1](#).

Monitor interfaces

Select to enable or disable monitoring FortiGate interfaces to verify the monitored interfaces are functioning properly and are connected to their networks. See [Link failover \(port monitoring or interface monitoring\) on page 1549](#).

If a monitored interface fails or is disconnected from its network, the interface leaves the cluster and a link failover occurs. The link failover causes the cluster to reroute the traffic being processed by that interface to the same interface of another cluster FortiGate that still has a connection to the network. This other cluster FortiGate becomes the new primary unit.

Interface monitoring (also called port monitoring) is disabled by default. Leave interface monitoring disabled until the cluster is operating and then only enable interface monitoring for connected interfaces.

You can monitor up to 64 interfaces.

Heartbeat interfaces

Enable or disable HA heartbeat communication for each interface in the cluster and set the heartbeat interface priority. The heartbeat interface with the highest priority processes all heartbeat traffic. If two or more heartbeat interfaces have the same priority, the heartbeat interface with the lowest hash map order value processes all heartbeat traffic. The GUI lists interfaces in alphanumeric order:

- port1
- port2 through 9
- port10

Hash map order sorts interfaces in the following order:

- port1
- port10
- port2 through port9

The default heartbeat interface configuration is different for each FortiGate model. This default configuration usually sets the priority of two heartbeat interfaces to 50. You can accept the default heartbeat interface configuration or change it as required.

The heartbeat interface priority range is 0 to 512. The default priority when you select a new heartbeat interface is 0.

You must select at least one heartbeat interface. If heartbeat communication is interrupted, the cluster stops processing traffic. See [HA heartbeat and communication between cluster units on page 1523](#).

You can select up to 8 heartbeat interfaces. This limit only applies to units with more than 8 physical interfaces.

Management interface reservation

You can provide direct management access to individual cluster units by reserving a management interface as part of the HA configuration. Once this management interface is reserved, you can configure a different IP address, administrative access and other interface settings for this interface for each cluster unit. You can also specify static routing settings for this interface. Then by connecting this interface of each cluster unit to your network you can manage each cluster unit separately from a different IP address. See [Managing individual cluster units using a reserved out-of-band management interface on page 1488](#).

VDOM partitioning

If you are configuring virtual clustering, you can set the virtual domains to be in virtual cluster 1 and the virtual domains to be in virtual cluster 2. The root virtual domain must always be in virtual cluster 1.

Secondary cluster settings

If you are configuring virtual clustering you can set the device priority and configure interface monitoring for the secondary virtual cluster.

FGCP configuration examples and troubleshooting

This chapter contains general procedures and descriptions as well as detailed configuration examples that describe how to configure FortiGate HA clusters. Some of the examples are available as cookbook recipes and this chapter provides introductions and links to the cookbook recipes.

About the examples in this chapter

The procedures in this chapter describe some of many possible sequences of steps for configuring HA clustering. As you become more experienced with FortiOS HA you may choose to use a different sequence of configuration steps.

For simplicity, many of these procedures assume that you are starting with new FortiGates set to the factory default configuration. However, starting from the default configuration is not a requirement for a successful HA deployment. FortiGate HA is flexible enough to support a successful configuration from many different starting points.

The examples in this chapter include example values only. In most cases you will substitute your own values. The examples in this chapter also do not contain detailed descriptions of configuration parameters.

How to set up FGCP clustering (recommended steps)

The following recipe describes how to enhance the reliability of a network protected by a FortiGate by adding a second FortiGate and setting up a FortiGate Clustering Protocol (FGCP) High Availability cluster.

High Availability with FGCP (Expert)

The following cookbook recipe also describes how to add a backup FortiGate to a previously installed FortiGate, to form a high availability (HA) cluster to improve network reliability. This recipe takes a higher-level approach using the GUI to configure HA.

High availability with two FortiGates

Adding a third FortiGate to an operating cluster and switching to active-active HA

This recipe describes how to add a third FortiGate to an already established FGCP cluster and how to configure active-active HA for that cluster.

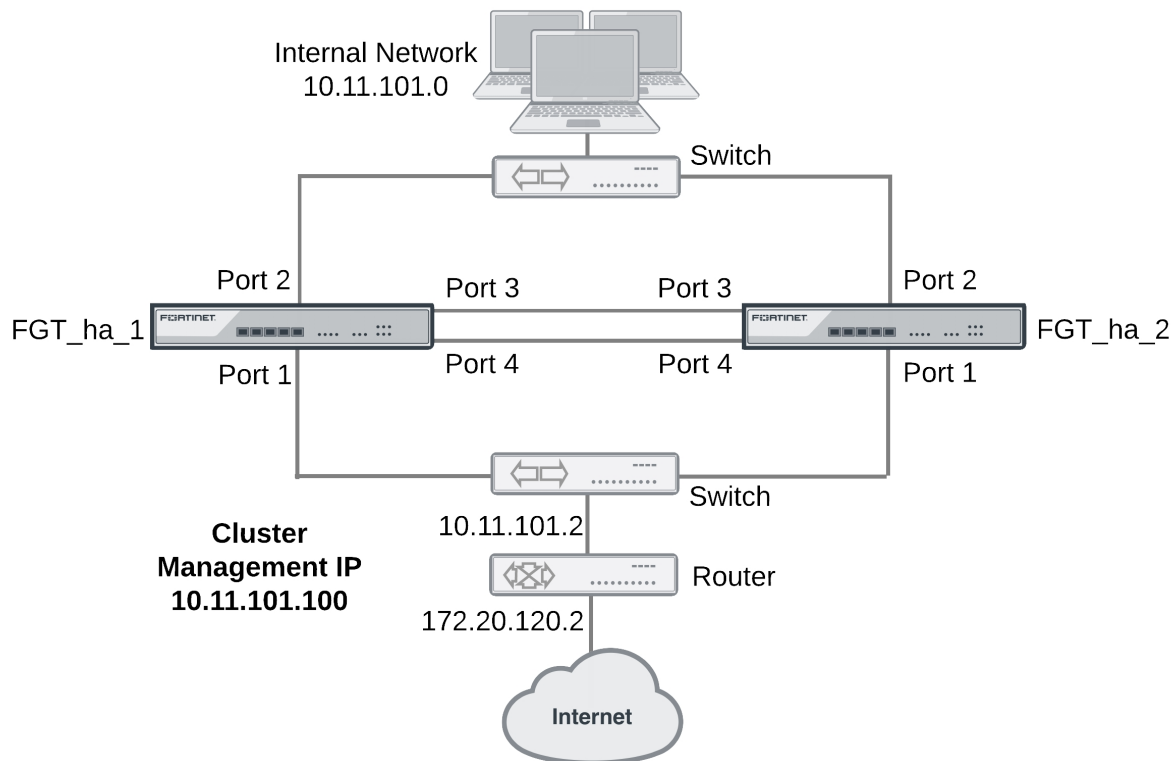
Adding a third FortiGate to an FGCP cluster

Active-active HA cluster in transparent mode

This section describes a simple HA network topology that includes an HA cluster of two generic FortiGates installed between an internal network and the internet and running in transparent mode.

Example transparent mode HA network topology

The figure below shows a transparent mode FortiGate HA cluster consisting of two FortiGates (FGT_ha_1 and FGT_ha_2) installed between the internet and internal network. The topology includes a router that performs NAT between the internal network and the internet. The cluster management IP address is 10.11.101.100.

transparent mode HA network topology

Port3 and port4 are used as the heartbeat interfaces. Because the cluster consists of two FortiGates, you can make the connections between the heartbeat interfaces using crossover cables. You could also use switches and regular ethernet cables.

General configuration steps

This section includes GUI and CLI procedures. These procedures assume that the FortiGates are running the same FortiOS firmware build and are set to the factory default configuration.

In this example, the configuration steps are identical to the NAT/Route mode configuration steps until the cluster is operating. When the cluster is operating, you can switch to transparent mode and add basic configuration settings to cluster.

General configuration steps

1. Apply licenses to the FortiGates to become the cluster.
2. Configure the FortiGates for HA operation.
 - Optionally change each unit's host name.
 - Configure HA.
2. Connect the cluster to the network.
3. Confirm that the cluster units are operating as a cluster.
4. Switch the cluster to transparent mode and add basic configuration settings to the cluster.

- Switch to transparent mode, add the management IP address and a default route.
- Add a password for the admin administrative account.
- View cluster status from the GUI or CLI.

Configuring a transparent mode active-active cluster of two FortiGates - GUI

Use the following procedures to configure the FortiGates for HA operation using the FortiGate GUI. These procedures assume you are starting with two FortiGates with factory default settings.



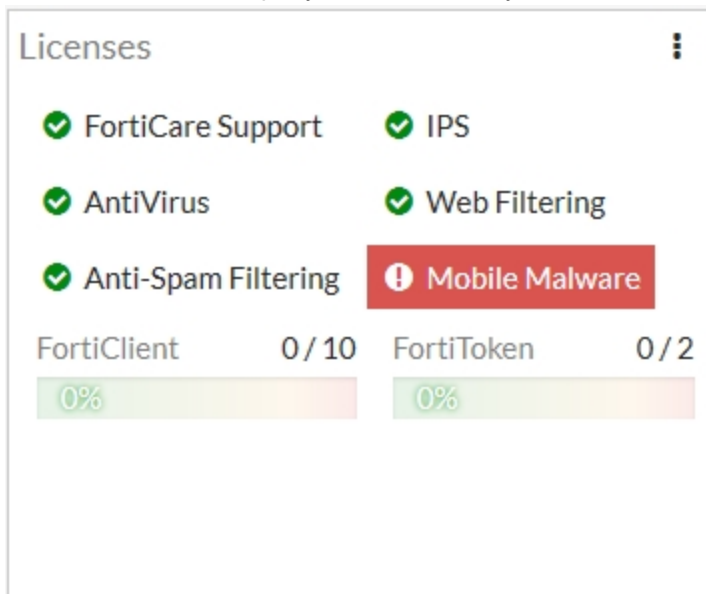
Waiting until you have established the cluster to switch to transparent mode means fewer configuration steps because you can switch the mode of the cluster in one step.

To configure the first FortiGate (host name FGT_ha_1)

1. Register and apply licenses to the FortiGate before configuring it for HA operation. This includes licensing for **FortiCare Support**, **IPS**, **AntiVirus**, **Web Filtering**, **Mobile Malware**, **FortiClient**, **FortiCloud**, and additional **virtual domains** (VDOMs). All FortiGates in the cluster must have the same level of licensing for FortiGuard, FortiCloud, FortiClient, and VDOMs. **FortiToken** licenses can be added at any time because they are synchronized to all cluster members.

If the FortiGates in the cluster will be running FortiOS Carrier, apply the FortiOS Carrier license before configuring the cluster (and before applying other licenses). Applying the FortiOS Carrier license sets the configuration to factory defaults, requiring you to repeat steps performed before applying the license.

You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed, third-party certificates are synchronized to the backup FortiGate.



2. Click on the System Information dashboard widget and select **Configure settings** in **System > Settings**.
3. Enter a new Host Name for this FortiGate.

New Name	FGT_ha_1
-----------------	----------

4. Select **OK**.
5. Go to **System > HA** and change the following settings:

Mode	Active-Active
Group Name	example2.com
Password	HA_pass_2



This is the minimum recommended configuration for an active-active HA cluster. You can configure other HA options at this point, but if you wait until the cluster is operating you will only have to configure these options once for the cluster instead of separately for each cluster unit.

6. Select **OK**.

The FortiGate negotiates to establish an HA cluster. When you select **OK** you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port2 interface virtual MAC: 00-09-0f-09-00-01
- port3 interface virtual MAC: 00-09-0f-09-00-02
- port4 interface virtual MAC: 00-09-0f-09-00-03

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

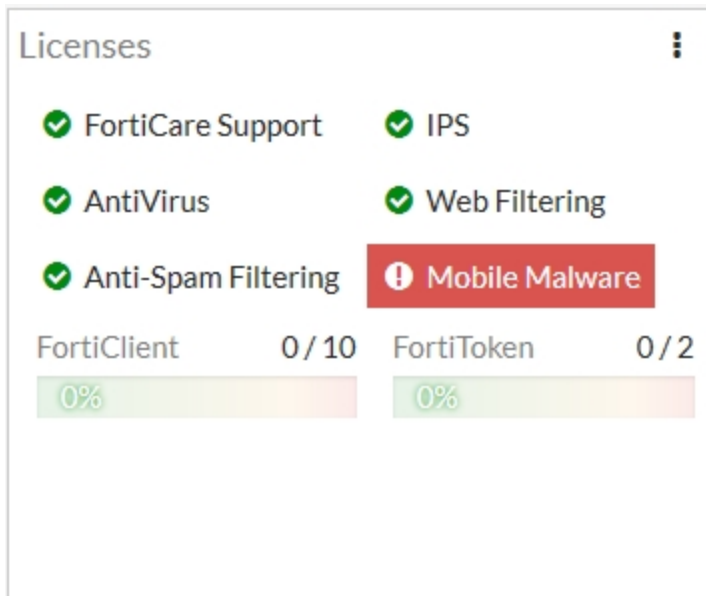
To confirm these MAC address changes, you can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (MAC) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
Current_HAaddr    00:09:0f:09:00:00
Permanent_HWaddr  02:09:0f:78:18:c9
.
.
.
```

10. Power off the first FortiGate.

To configure the second FortiGate (host name FGT_ha_2)

1. Register and apply licenses to the FortiGate before configuring it for HA operation.



2. Click on the System Information dashboard widget and select **Configure settings** in **System > Settings**.
3. Enter a new Host Name for this FortiGate.

New Name	FGT_ha_2
-----------------	----------

4. Select **OK**.
5. Go to **System > HA** and change the following settings:

Mode	Active-Active
Group Name	example2.com
Password	HA_pass_2

6. Select **OK**.

The FortiGate negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate interfaces.

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

7. Power off the second FortiGate.

To connect the cluster to the network

1. Connect the port1 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internet.
2. Connect the port2 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internal network.

3. Connect the port3 interfaces of FGT_ha_1 and FGT_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
5. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

To switch the cluster to transparent mode

Switching from NAT/Route to transparent mode involves adding the transparent mode management IP address and default route.



This is the minimum recommended configuration for an active-active HA cluster. You can configure other HA options at this point, but if you wait until the cluster is operating you will only have to configure these options once for the cluster instead of separately for each cluster unit.

1. Start a web browser and browse to the address `https://192.168.1.99` (remember to include the “s” in `https://`).
The FortiGate Login is displayed.
2. Type admin in the Name field and select Login.
3. Under System Information, beside **Operation Mode** select **Change**.
4. Set Operation Mode to transparent.
5. Configure basic transparent mode settings.

Operation Mode	Transparent
Management IP/Mask	10.11.101.100/24
Default Gateway	10.11.101.2

6. Select **Apply**.
The cluster switches to operating in transparent mode. The virtual MAC addresses assigned to the cluster interfaces do not change.

To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.



Once the cluster is operating, because configuration changes are synchronized to all cluster units, configuring the cluster is the same as configuring an individual FortiGate. You could have performed the following configuration steps separately on each FortiGate before you connected them to form a cluster.

1. Start a web browser and browse to the address `https://10.11.101.100` (remember to include the “s” in `https://`).

The FortiGate Login is displayed.

2. Type **admin** in the Name field and select **Login**.

The FortiGate dashboard is displayed.

The HA Status dashboard widget displays how long the cluster has been operating (Uptime) and the time since the last failover occurred (State Changed). You can hover over the State Changed time to see the event that caused the state change. You can also click on the HA Status dashboard widget to configure HA settings or to get a listing of the most recent HA events recorded by the cluster.

3. Go to **System > HA** to view the cluster members list.

The list shows both cluster units, their host names, their roles in the cluster, and their device priorities. You can use this list to confirm that the cluster is operating normally. For example, if the list shows only one cluster unit then the other unit has left the cluster for some reason.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units, the FortiGates are not functioning as a cluster. See [Troubleshooting HA clusters on page 1458](#) to troubleshoot the cluster.

To add basic configuration settings to the cluster

Use the following steps to configure the cluster. Note that the following are example configuration steps only and do not represent all of the steps required to configure the cluster for a given network.

1. Log into the cluster GUI.
2. Go to **System > Administrators**.
3. Edit **admin** and select **Change Password**.
4. Enter and confirm a new password.
5. Select **OK**.



You added a default gateway when you switched to transparent mode so you don't need to add a default route as part of the basic configuration of the cluster at this point.

Configuring a transparent mode active-active cluster of two FortiGates - CLI

Use the following procedures to configure the FortiGates for transparent mode HA operation using the FortiGate CLI.

To configure each FortiGate for HA operation

1. Power on the FortiGate.
2. Connect a null modem cable to the communications port of the management computer and to the FortiGate Console port.
3. Start HyperTerminal, enter a name for the connection, and select **OK**.
4. Configure HyperTerminal to connect directly to the communications port on the computer to which you have connected the null modem cable and select **OK**.
5. Select the following port settings and select **OK**.

Bits per second	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	None

6. Press **Enter** to connect to the FortiGate CLI.

The FortiGate CLI login prompt appears. If the prompt does not appear, press Enter. If it still does not appear, power off your FortiGate and power it back on. If you are connected, at this stage you will see startup messages that will confirm you are connected. The login prompt will appear after the startup has completed.

7. Type `admin` and press **Enter** twice.

8. Register and apply licenses to the FortiGate.

9. Change the host name for this FortiGate. For example:

```
config system global
  set hostname FGT_ha_1
end
```

10. Configure HA settings.

```
config system ha
  set mode a-a
  set group-name example2.com
  set password HA_pass_2
end
```



This is the minimum recommended configuration for an active-active HA cluster. You can also configure other HA options, but if you wait until after the cluster is operating you will only have to configure these options once for the cluster instead of separately for each cluster unit.

The FortiGate negotiates to establish an HA cluster. You may temporarily lose network connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port2 interface virtual MAC: 00-09-0f-09-00-01
- port3 interface virtual MAC: 00-09-0f-09-00-02
- port4 interface virtual MAC: 00-09-0f-09-00-03

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

To confirm these MAC address changes, you can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate

interface. For example, use the following command to view the port1 interface virtual MAC address (MAC) and the port1 permanent MAC address (Permanent_HWaddr):

```
get hardware nic port1
.
.
.
Current_HAaddr    00:09:0f:09:00:00
Permanent_HWaddr  02:09:0f:78:18:c9
.
.
.
```

10. Display the HA configuration (optional).

```
get system ha
group-id : 0
group-name : example2.com
mode : a-a
password : *
hbdev : "port3" 50 "port4" 50
session-sync-dev :
route-ttl : 10
route-wait : 0
route-hold : 10
sync-config : enable
encryption : disable
authentication : disable
hb-interval : 2
hb-lost-threshold : 20
hello-holddown : 20
arps : 5
arps-interval : 8
session-pickup : disable
update-all-session-timer: disable
session-sync-daemon-number: 1
link-failed-signal : disable
uninterruptible-upgrade: enable
ha-mgmt-status : disable
ha-eth-type : 8890
hc-eth-type : 8891
l2ep-eth-type : 8893
ha-uptime-diff-margin: 300
vcluster2 : disable
vcluster-id : 1
override : disable
priority : 128
slave-switch-standby: disable
minimum-worker-threshold: 1
monitor :
pingserver-monitor-interface:
pingserver-failover-threshold: 0
pingserver-slave-force-reset: enable
pingserver-flip-timeout: 60
vdom : "root"
```

11. Power off the FortiGate.

To configure the second FortiGate (host name FGT_ha_2)

1. Power on the FortiGate.
2. Connect a null modem cable to the communications port of the management computer and to the FortiGate Console port.
3. Start HyperTerminal, enter a name for the connection, and select **OK**.
4. Configure HyperTerminal to connect directly to the communications port on the computer to which you have connected the null modem cable and select **OK**.
5. Select the following port settings and select **OK**.

Bits per second	9600
Data bits	8
Parity	None
Stop bits	1
Flow control	None

6. Press **Enter** to connect to the FortiGate CLI.
The FortiGate CLI login prompt appears. If the prompt does not appear, press Enter. If it still does not appear, power off your FortiGate and power it back on. If you are connected, at this stage you will see startup messages that will confirm you are connected. The login prompt will appear after the startup has completed.
7. Type `admin` and press **Enter** twice.
8. Register and apply licenses to the FortiGate.
9. Change the host name for this FortiGate.

```
config system global
    set hostname FGT_ha_2
end
```

10. Configure HA settings.

```
config system ha
    set mode a-a
    set group-name example2.com
    set password HA_pass_2
end
```

The FortiGate negotiates to establish an HA cluster. You may temporarily lose network connectivity with the FortiGate as the HA cluster negotiates and because the FGCP changes the MAC address of the FortiGate interfaces.

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

11. Display the HA configuration (optional).

```
get system ha
group-id : 0
group-name : example2.com
mode : a-a
password : *
```

```

hbdev : "port3" 50 "port4" 50
session-sync-dev :
route-ttl : 10
route-wait : 0
route-hold : 10
sync-config : enable
encryption : disable
authentication : disable
hb-interval : 2
hb-lost-threshold : 20
hello-holddown : 20
arps : 5
arps-interval : 8
session-pickup : disable
update-all-session-timer: disable
session-sync-daemon-number: 1
link-failed-signal : disable
uninterruptible-upgrade: enable
ha-mgmt-status : disable
ha-eth-type : 8890
hc-eth-type : 8891
l2ep-eth-type : 8893
ha-uptime-diff-margin: 300
vcluster2 : disable
vcluster-id : 1
override : disable
priority : 128
schedule : round-robin
monitor :
pingserver-monitor-interface:
pingserver-failover-threshold: 0
pingserver-slave-force-reset: enable
pingserver-flip-timeout: 60
vdom : "root"
schedule : round-robin

```

12. Power off the FortiGate.

To connect the cluster to the network

1. Connect the port1 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internet.
2. Connect the port2 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internal network.
3. Connect the port3 interfaces of FGT_ha_1 and FGT_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port4 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
5. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

To connect to the cluster CLI and switch the cluster to transparent mode

1. Determine which cluster unit is the primary unit.

- Use the null-modem cable and serial connection to re-connect to the CLI of one of the cluster units.
- Enter the command `get system status`.
- If the command output includes `Current HA mode: a-a, master`, the cluster units are operating as a cluster and you have connected to the primary unit. Continue with Step 2.
- If the command output includes `Current HA mode: a-a, backup`, you have connected to a subordinate unit. Connect to the other cluster unit, which should be the primary unit and continue with Step 2.



If the command output includes `Current HA mode: standalone`, the cluster unit is not operating in HA mode.

2. Change to transparent mode.

```
config system settings
  set opmode transparent
  set manageip 192.168.20.3/24
  set gateway 192.168.20.1
end
```

The cluster switches to transparent Mode, and your administration session is disconnected.

You can now connect to the cluster CLI using SSH to connect to the cluster internal interface using the management IP address (192.168.20.3).

To view cluster status

Use the following steps to view cluster status from the CLI.

1. Determine which cluster unit is the primary unit.

- Use the null-modem cable and serial connection to re-connect to the CLI of one of the cluster units.
- Enter the command `get system status`.
- If the command output includes `Current HA mode: a-a, master`, the cluster units are operating as a cluster and you have connected to the primary unit. Continue with Step ["Active-active HA cluster in transparent mode" on page 1412](#).
- If the command output includes `Current HA mode: a-a, backup`, you have connected to a subordinate unit. Connect the null-modem cable to the other cluster unit, which should be the primary unit and continue with Step 2.



If the command output includes `Current HA mode: standalone`, the cluster unit is not operating in HA mode and you should review your HA configuration.

2. Enter the following command to confirm the HA configuration of the cluster:

```
get system ha status
HA Health Status: OK
Model: FortiGate-XXXX
Mode: HA A-P
Group: 0
Debug: 0
Cluster Uptime: 7 days 00:30:26
.
```

.

.

You can use this command to confirm that the cluster is healthy and operating normally, some information about the cluster configuration, and information about how long the cluster has been operating. Information not shown in this example includes how the primary unit was selected, configuration synchronization status, usage stats for each cluster unit, heartbeat status, and the relative priorities of the cluster units.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard do not display information for both cluster units the FortiGates are not functioning as a cluster. See [Troubleshooting HA clusters on page 1458](#) to troubleshoot the cluster.

To add a password for the admin administrative account

1. Add a password for the admin administrative account.

```
config system admin
  edit admin
    set password <psswr>
end
```

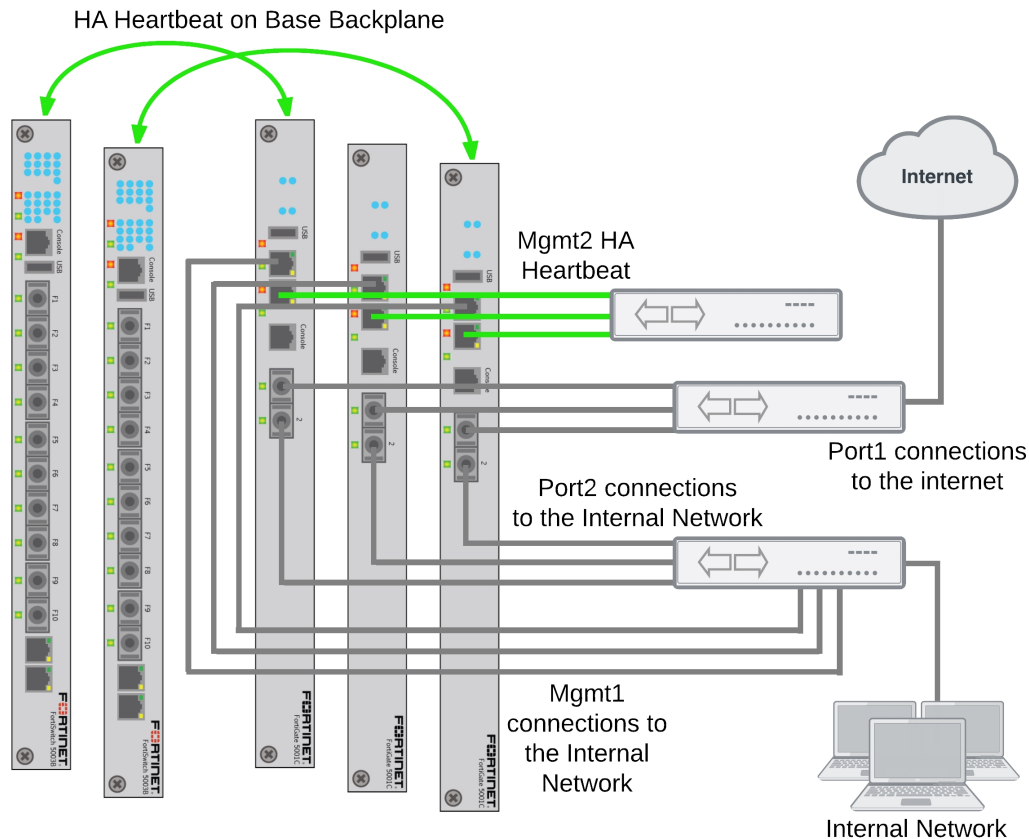
FortiGate-5000 active-active HA cluster with FortiClient licenses

This section describes how to configure an HA cluster of three FortiGate-5001D units that connect an internal network to the internet. The FortiGate-5001D units each have a FortiClient license installed on them to support FortiClient profiles.

Normally it is recommended that you add FortiClient licenses to the FortiGates before setting up the cluster. This example; however, describes how to apply FortiClient licenses to the FortiGates in an operating cluster.

Example network topology

The following diagram shows an HA cluster consisting of three FortiGate-5001D cluster units (host names slot-3, slot-4, and slot-5) installed in a FortiGate-5000 series chassis with two FortiController-5003B units for heartbeat communication between the cluster units. The cluster applies security features including FortiClient profiles to data traffic passing through it.



The cluster is managed from the internal network using the FortiGate-5001D mgmt1 interfaces configured as HA reserved management interfaces. Using these reserved management interfaces the overall cluster can be managed and cluster units can be managed individually. Individual management access to each cluster unit makes some operations, such as installing FortiClient licenses, easier and also allows you to view status of each cluster unit.

The reserved management interface of each cluster unit has a different IP address and retains its own MAC address. The cluster does not change the reserved management interface MAC address.

Example network topology

By default base1 and base2 are used for heartbeat communication between the FortiGates. To use the base1 and base2 interfaces for the HA heartbeat, the example describes how to display the backplane interfaces on the GUI before turning on HA.

This example also includes using the mgmt2 interface for heartbeat communication for additional heartbeat redundancy.

To connect the cluster

1. Connect the FortiGate-5001D port1 interfaces to a switch and connect that switch to the internet.
2. Connect the FortiGate-5001D port2 interfaces to a switch and connect that switch to the internal network.
3. Connect the FortiGate-5001D mgmt1 interfaces to a switch that connects to the engineering network.
4. Connect the FortiGate-5001D mgmt2 interfaces to a switch for heartbeat communication between them.

Configuring the FortiGate-5000 active-active cluster - GUI

These procedures assume you are starting with three FortiGate-5001D boards and two FortiSwitch-5003B boards installed in a compatible FortiGate-5000 series chassis. The FortiSwitch-5003B boards are in chassis slots 1 and 2 and the FortiGate-5001D boards are in chassis slots 3, 4, and 5 and the chassis is powered on. All devices are in their factory default configuration. No configuration changes to the FortiSwitch-5003B boards are required.

To configure the FortiGate-5001D units

1. From the internal network, log into the GUI of the FortiGate-5001D unit in chassis slot 3 by connecting to the mgmt1 interface.

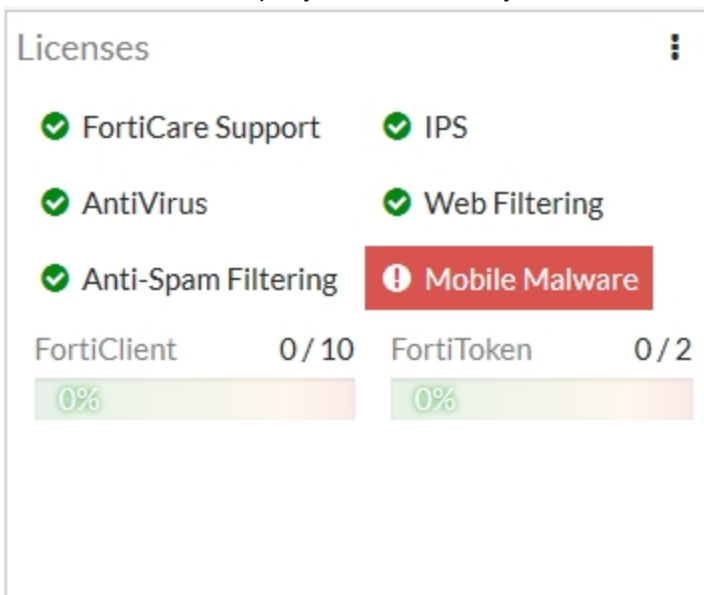


By default the mgmt1 interface of each FortiGate-5001D unit has the same IP address. To log into each FortiGate-5001D unit separately you could either disconnect the mgmt1 interfaces of the units that you don't want to log into or change the mgmt1 interface IP addresses for each unit by connecting to each unit's CLI from their console port.

2. Register and apply licenses to the FortiGate before configuring it for HA operation. This includes licensing for **FortiCare Support**, **IPS**, **AntiVirus**, **Web Filtering**, **Mobile Malware**, **FortiCloud**, and additional **virtual domains** (VDOMs). All FortiGates in the cluster must have the same level of licensing for FortiGuard, FortiCloud, FortiClient, and VDOMs. **FortiToken** licenses can be added at any time because they are synchronized to all cluster members. FortiClient licenses will be added in a following step.

If the FortiGates in the cluster will be running FortiOS Carrier, apply the FortiOS Carrier license before configuring the cluster (and before applying other licenses). Applying the FortiOS Carrier license sets the configuration to factory defaults, requiring you to repeat steps performed before applying the license.

You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed, third-party certificates are synchronized to the backup FortiGate.



3. Click on the System Information dashboard widget and select **Configure settings** in **System > Settings**.
4. Enter a new Host Name for this FortiGate, for example:

New Name	5001D-Slot-3
-----------------	--------------

5. Connect to the CLI and enter the following command to display backplane interfaces on the GUI:

```
config system global
  set show-backplane-intf enable
end
```

6. Set the Administrative Status of the base1 and base 2 interfaces to **Up**.

You can do this from the GUI by going to **Network > Interfaces**, editing each interface and setting **Administrative Status** to **Up**.

You can also do this from the CLI using the following command:

```
config system interface
  edit base1
    set status up
  next
  edit base2
    set status up
end
```

7. Go to **Network > Interfaces** and configure the IP address of the mgmt1 interface.

Because mgmt1 will become the reserved management interface for the cluster unit each FortiGate-5001D should have a different mgmt1 interface IP address. Give the mgmt1 interface an address that is valid for the internal network. Once HA with the reserved Management interface is enabled the IP address of the mgmt1 interface can be on the same subnet as the port2 interface (which will also be connected to the Internal network).

After the FortiGate is operating in HA mode the mgmt1 interface will retain its original MAC address instead of being assigned a virtual MAC address.

8. Go to **System > HA** and change the following settings:

Set the **Mode** to **Active-Active**.

Select **Reserve Management Port for Cluster Member** and select **mgmt1**.

Set the group name and password:

Group Name	example3.com
Password	HA_pass_3

Set the Heartbeat interface configuration to use base1, base2 and mgmt2 for heartbeat communication. Set the priority of each heartbeat interface to 50:

Heartbeat Interface		
	Enable	Priority
base1	Select	50
base2	Select	50
mgmt2	Select	50

9. Select **OK**.

The FortiGate negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate-5001D interfaces change to the following virtual MAC addresses:

- base1 interface virtual MAC: 00-09-0f-09-00-00
- base2 interface virtual MAC: 00-09-0f-09-00-01
- fabric1 interface virtual MAC: 00-09-0f-09-00-02
- fabric2 interface virtual MAC: 00-09-0f-09-00-03
- fabric3 interface virtual MAC: 00-09-0f-09-00-04
- fabric4 interface virtual MAC: 00-09-0f-09-00-05
- fabric5 interface virtual MAC: 00-09-0f-09-00-06
- mgmt1 keeps its original MAC address
- mgmt2 interface virtual MAC: 00-09-0f-09-00-08
- port1 interface virtual MAC: 00-09-0f-09-00-09
- port2 interface virtual MAC: 00-09-0f-09-00-0a

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic base1
.
.
.
Current_HWaddr 00:09:0f:09:00:00
Permanent_HWaddr 00:09:0f:71:0a:dc
.
.
.
```

9. Repeat these steps for the FortiGate-5001D units in chassis slots 4 and 5, with the following differences. Set the mgmt1 interface IP address of each FortiGate-5001D unit to a different IP address.

Set the FortiGate-5001D unit in chassis slot 4 host name to:

New Name	5001D-Slot-4
-----------------	--------------

Set the FortiGate-5001D unit in chassis slot 5 host name to:

New Name	5001D-Slot-5
-----------------	--------------

As you configure each FortiGate, they will negotiate and join the cluster.

To view cluster status

As you add units to the cluster you can log into the GUI of one of the cluster units to view the status of the cluster. The status displays will show each unit as it is added to the cluster.

1. Log into the primary unit or any cluster unit and view the system dashboard.
The HA Status dashboard widget displays how long the cluster has been operating (Uptime) and the time since the last failover occurred (State Changed) You can hover over the State Changed time to see the event that caused the state change You can also click on the HA Status dashboard widget to configure HA settings or to get a listing of the most recent HA events recorded by the cluster.
2. Go to **System > HA** to view the cluster members list.
The list shows both cluster units, their host names, their roles in the cluster, and their device priorities. You can use this list to confirm that the cluster is operating normally. For example, if the list shows only one cluster unit then the other unit has left the cluster for some reason.

To troubleshoot the cluster

See [Troubleshooting HA clusters on page 1458](#).

To manage each cluster unit

Because you have configured a reserved management interface, you can manage each cluster unit separately by connecting to the IP address you configured for each unit's mgmt1 interface. You can view the status of each cluster unit and make changes to each unit's configuration. For example, as described below, each cluster unit must have its own FortiClient license. You can use the reserved management IP addresses to connect to each cluster unit to install the FortiClient license for that unit.

Usually you would make configuration changes by connecting to the primary unit and changing its configuration. The cluster then synchronizes the configuration changes to all cluster units. If you connect to individual cluster units and change their configuration, those configuration changes are also synchronized to each cluster unit. The exception to this is configuration objects that are not synchronized, such as the host name, FortiClient license and so on.

You can also manage each cluster unit by logging into the primary unit CLI and using the following command to connect to other cluster units:

```
execute ha manage <cluster-index>
```

To add basic configuration settings to the cluster

Use the following steps to configure the cluster.

1. Log into the cluster GUI.
You can log into the primary unit or any one of the cluster units using the appropriate mgmt1 IP address.
2. Go to **System > Administrators**.
3. Edit **admin** and select **Change Password**.
4. Enter and confirm a new password.
5. Select **OK**.
6. Go to **Network > Interfaces** and edit the **port1** interface. Set this interface IP address to the address required to connect to the interface to the internet.
7. Edit the port2 interface and set its IP to an IP address for the internal network.

To add a FortiClient license to each cluster unit

Normally you would add FortiClient licenses to the FortiGates before forming the cluster. However, you can use the following steps to add FortiClient licenses to an operating cluster.

Contact your reseller to purchase FortiClient licenses for your cluster units. Each cluster unit must have its own FortiClient license.

When you receive the license keys you can visit the [Fortinet Support](#) website and add a FortiClient license key to each licensed FortiGate. Then, as long as the cluster can connect to the internet the license keys are downloaded from the FortiGuard network to all of the FortiGates in the cluster.

You can also use the following steps to manually add the license keys to your cluster units from the GUI. Your cluster must be connected to the internet.

1. Log into the GUI of each cluster unit using its reserved management interface IP address.
2. Go to the **License Information** dashboard widget and beside FortiClient select **Enter License**.
3. Enter the license key and select **OK**.
4. Confirm that the license has been installed and the correct number of FortiClients are licensed.
5. Repeat for all of the cluster units.

You can also use the following command to add the license key from the CLI:

```
execute FortiClient-NAC update-registration-license <license-number>
```

You can connect to the CLIs of each cluster unit using their reserved management IP address.

You can also log into the primary unit CLI and use the `execute ha manage` command to connect to each cluster unit CLI.

Configuring the FortiGate-5000 active-active cluster - CLI

These procedures assume you are starting with three FortiGate-5001D boards and two FortiSwitch-5003B boards installed in a compatible FortiGate-5000 series chassis. The FortiSwitch-5003B boards are in chassis slots 1 and 2 and the FortiGate-5001D boards are in chassis slots 3, 4, and 5 and the chassis is powered on. All devices are in their factory default configuration. No configuration changes to the FortiSwitch-5003B boards are required.

To configure the FortiGate-5005FA2 units

1. From the internal network, log into the CLI of the FortiGate-5001D unit in chassis slot 3 by connecting to the mgmt1 interface.



By default the mgmt1 interface of each FortiGate-5001D unit has the same IP address. To log into each FortiGate-5001D unit separately you could either disconnect the mgmt1 interfaces of the units that you don't want to log into or change the mgmt1 interface IP addresses for each unit by connecting to each unit's CLI from their console port.

You can also use a console connection.

2. Register and apply licenses to the FortiGate.
3. Change the host name for this FortiGate. For example:

```
config system global
  set hostname 5001D-Slot-3
end
```

4. Enter the following command to display backplane interfaces on the GUI:

```
config system global
  set show-backplane-intf enable
end
```

5. Set the Administrative Status of the base1 and base 2 interfaces to **Up**.

```
config system interface
  edit base1
    set status up
  next
  edit base2
    set status up
end
```

6. Add an IP address to the mgmt1 interface.

```
config system interface
  edit mgmt1
    set ip 172.20.120.110/24
    set allowaccess http https ssl ping
  end
```

Because mgmt1 will become the reserved management interface for the cluster unit each FortiGate-5001D should have a different mgmt1 interface IP address. Give the mgmt1 interface an address that is valid for the internal network. Once HA with the reserved Management interface is enabled the IP address of the mgmt1 interface can be on the same subnet as the port2 interface (which will also be connected to the Internal network).

7. Configure HA settings.

```
config system ha
  set mode a-a
  set ha-mgmt-status enable
  set ha-mgmt-interface mgmt1
  set group-name example3.com
  set password HA_pass_3
  set hbdev base1 50 base2 50 mgmt2 50
end
```

The FortiGate negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate-5001D interfaces change to the following virtual MAC addresses:

- base1 interface virtual MAC: 00-09-0f-09-00-00
- base2 interface virtual MAC: 00-09-0f-09-00-01
- fabric1 interface virtual MAC: 00-09-0f-09-00-02
- fabric2 interface virtual MAC: 00-09-0f-09-00-03
- fabric3 interface virtual MAC: 00-09-0f-09-00-04
- fabric4 interface virtual MAC: 00-09-0f-09-00-05
- fabric5 interface virtual MAC: 00-09-0f-09-00-06
- mgmt1 keeps its original MAC address
- mgmt2 interface virtual MAC: 00-09-0f-09-00-08
- port1 interface virtual MAC: 00-09-0f-09-00-09
- port2 interface virtual MAC: 00-09-0f-09-00-0a

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic base1
.
.
.
Current_HWaddr 00:09:0f:09:00:00
Permanent_HWaddr 00:09:0f:71:0a:dc
.
.
.
```

7. Repeat these steps for the FortiGate-5001D units in chassis slots 4 and 5, with the following differences. Set the mgmt1 interface IP address of each FortiGate-5001D unit to a different IP address.

Set the FortiGate-5001D unit in chassis slot 4 host name to:

```
config system global
    set hostname 5001D-Slot-4
end
```

Set the FortiGate-5001D unit in chassis slot 5 host name to:

```
config system global
    set hostname 5001D-Slot-5
end
```

As you configure each FortiGate, they will negotiate and join the cluster.

To view cluster status

As you add units to the cluster you can log into the CLI of one of the cluster units using its reserved management interface to view the status of the cluster. The status will show each unit as it is added to the cluster.

For example, the following command output shows the status of the cluster when all three cluster units have been added:

```
get system ha status
HA Health Status: OK
Model: FortiGate-XXXX
Mode: HA A-P
Group: 0
Debug: 0
Cluster Uptime: 7 days 00:30:26
.
.
.
Slave : 5001d-slot4      , FG-5KD3914800284, operating cluster index = 2
Master: 5001d-slot5     , FG-5KD3914800353, operating cluster index = 0
Slave : 5001d-slot3     , FG-5KD3914800344, operating cluster index = 1
```

You can use this command to confirm that the cluster is healthy and operating normally, some information about the cluster configuration, and information about how long the cluster has been operating. Information not shown in this example includes how the primary unit was selected, configuration synchronization status, usage stats for each cluster unit, heartbeat status, and the relative priorities of the cluster units.

To troubleshoot the cluster

See [Troubleshooting HA clusters on page 1458](#).

To manage each cluster unit

Because you have configured a reserved management interface, you can manage each cluster unit separately by connecting to the IP address you configured for each unit's mgmt1 interface. You can view the status of each cluster unit and make changes to each unit's configuration. For example, as described below, each cluster unit must have its own FortiClient license. You can use the reserved management IP addresses to connect to each cluster unit to install the FortiClient license for that unit.

Usually you would make configuration changes by connecting to the primary unit and changing its configuration. The cluster then synchronizes the configuration changes to all cluster units. If you connect to individual cluster units and change their configuration, those configuration changes are also synchronized to each cluster unit. The exception to this is configuration objects that are not synchronized, such as the host name, FortiClient license and so on.

You can also manage each cluster unit by logging into the primary unit CLI and using the following command to connect to other cluster units:

```
execute ha manage <cluster-index>
```

To add a password for the admin administrative account

1. Add a password for the admin administrative account.

```
config system admin
  edit admin
    set password <psswr>
  end
```

To add basic configuration settings to the cluster

Use the following steps to configure the cluster.

1. Log into the cluster CLI.

You can log into the primary unit or any one of the cluster units using the appropriate mgmt1 IP address.

2. Add a password for the admin administrative account.

```
config system admin
  edit admin
    set password <psswr>
  end
```

3. Set the port1 interface IP address to the address required to connect to the interface to the internet.

```
config system interface
  edit port1
    set ip 10.10.10.10/24
  end
```


4. Set the port2 interface IP address to the address required to connect to the interface to the internal network.

```
config system interface
edit port2
set ip 172.20.120.12/24
end
```

To add a FortiClient license to each cluster unit

Normally you would add FortiClient licenses to the FortiGates before forming the cluster. However, you can use the following steps to add FortiClient licenses to an operating cluster.

Contact your reseller to purchase FortiClient licenses for your cluster units. Each cluster unit must have its own FortiClient license.

When you receive the license keys you can visit the [Fortinet Support](#) website and add a FortiClient license key to each licensed FortiGate. Then, as long as the cluster can connect to the internet the license keys are downloaded from the FortiGuard network to all of the FortiGates in the cluster.

You can also use the following steps to manually add the license keys to your cluster units from the CLI. Your cluster must be connected to the internet.

1. Log into the CLI of each cluster unit using its reserved management interface IP address.
2. Enter the following command to the unit's serial number:

```
get system status
```

3. Enter the following command to add the license key for that serial number:

```
execute FortiClient-NAC update-registration-license <license-key>
```

4. Confirm that the license has been installed and the correct number of FortiClients are licensed.

```
execute forticlient info
Maximum FortiClient connections: unlimited.
Licensed connections: 114
    NAC: 114
    WANOPT: 0
    Test: 0
Other connections:
    IPsec: 0
    SSLVPN: 0
```

5. Repeat for all of the cluster units.

You can also log into the primary unit CLI and use the `execute ha manage` command to connect to each cluster unit CLI.

Replacing a failed cluster unit

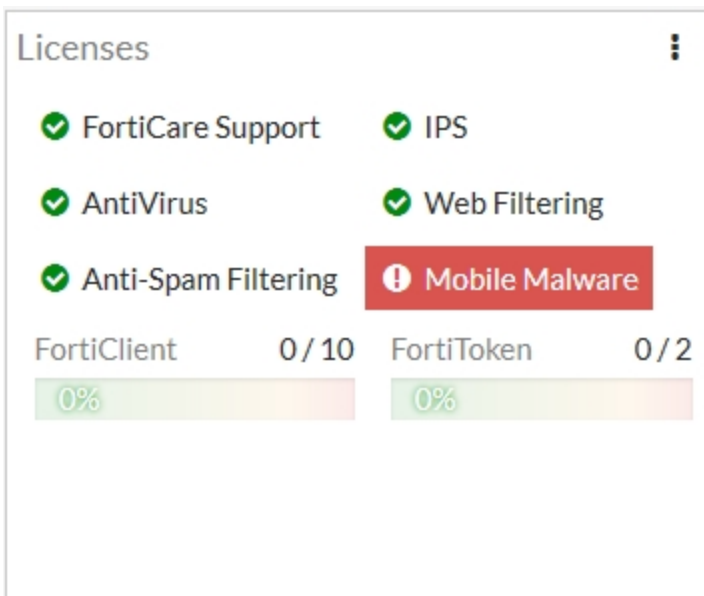
This procedure describes how to remove a failed cluster unit from a cluster and add a new one to replace it. You can also use this procedure to remove a failed unit from a cluster, repair it and add it back to the cluster.

Replacing a failed does not interrupt the operation of the cluster unless you have to change how the cluster is connected to the network to accommodate the replacement unit.

You can use this procedure to replace more than one cluster unit.

To replace a failed cluster unit

1. Disconnect the failed unit from the cluster and the network.
If you maintain other connections between the network and the still functioning cluster unit or units and between remaining cluster units network traffic will continue to be processed.
2. Repair the failed cluster unit, or obtain a replacement unit with the exact same hardware configuration as the failed cluster unit.
3. Install the same firmware build on the repaired or replacement unit as is running on the cluster.
4. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS). All of the FortiGates in a cluster must have the same level of licensing.



5. You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate.
We recommend that you add FortiToken licenses and FortiTokens to the primary unit after the cluster has formed.
6. Configure the repaired or replacement unit for HA operation with the same HA configuration as the cluster.
7. If the cluster is running in transparent mode, change the operating mode of the repaired or replacement unit to transparent mode.
8. Connect the repaired or replacement cluster unit to the cluster.
For an example see [How to set up FGCP clustering \(recommended steps\) on page 1412](#).
9. Power on the repaired or replacement cluster unit.
When the unit starts it negotiates to join the cluster. After it joins the cluster, the cluster synchronizes the repaired or replacement unit configuration with the configuration of the primary unit.

You can add a repaired or replacement unit to a functioning cluster at any time. The repaired or replacement cluster unit must:

- Have the same hardware configuration as the cluster units. Including the same hard disk configuration and the same AMC cards installed in the same slots.
- Have the same firmware build as the cluster.

- Be set to the same operating mode (NAT or transparent) as the cluster.
- Be operating in single VDOM mode.

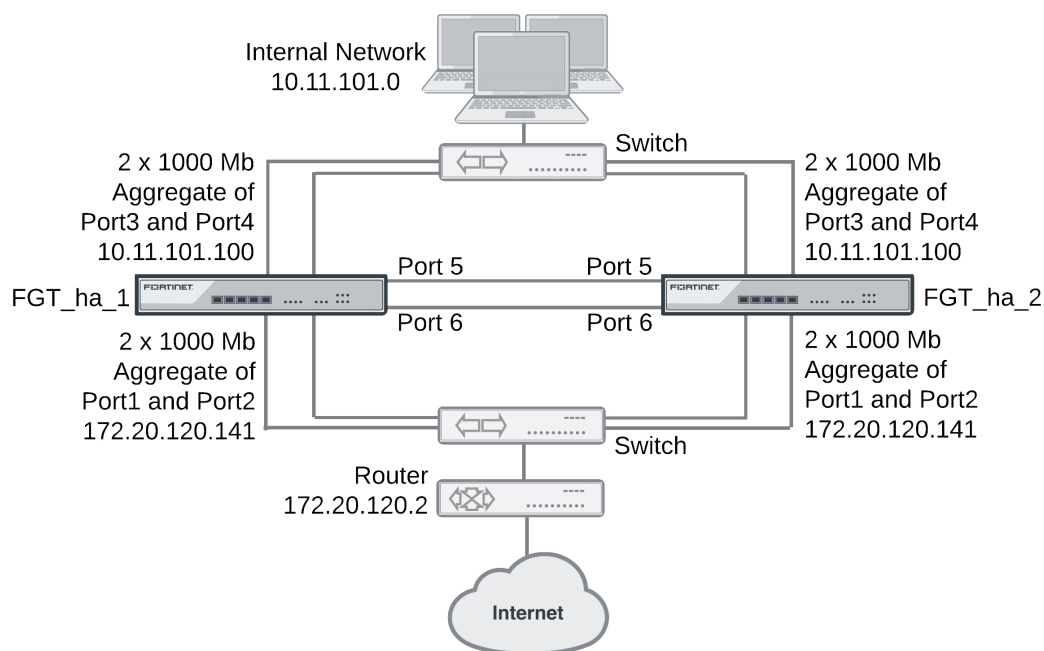
FGCP HA with 802.3ad aggregated interfaces

On FortiGate models that support it you can use 802.3ad link aggregation to combine two or more interfaces into a single aggregated interface. 802.3ad Link Aggregation and its management protocol, Link Aggregation Control Protocol (LACP) are a method for combining multiple physical links into a single logical link. This increases both potential throughput and network resiliency. Using LACP, traffic is distributed among the physical interfaces in the link, potentially resulting in increased performance.

This example describes how to configure an HA cluster consisting of two FortiGates with two aggregated 1000 Mb connections to the internet using port1 and port2 and two aggregated 1000 Mb connections to the internal network using port3 and port4. The aggregated interfaces are also configured as HA monitored interfaces.

Each of the aggregate links connects to a different switch. Each switch is configured for link aggregation (2x1000Mb).

Example cluster with aggregate interfaces



HA interface monitoring, link failover, and 802.3ad aggregation

When monitoring the aggregated interface, HA interface monitoring treats the aggregated link as a single interface and does not monitor the individual physical interfaces in the link. HA interface monitoring registers the link to have failed only if all the physical interfaces in the link have failed. If only some of the physical interfaces in the link fail or become disconnected, HA considers the link to be operating normally.

HA MAC addresses and 802.3ad aggregation

If a configuration uses the Link Aggregate Control Protocol (LACP) (either passive or active), LACP is negotiated over all of the interfaces in any link. For a standalone FortiGate, the FortiGate LACP implementation uses the MAC address of the first interface in the link to uniquely identify that link. For example, a link consisting of port1 and port2 interfaces would have the MAC address of port1.

In an HA cluster, HA changes the MAC addresses of the cluster interfaces to virtual MAC addresses. An aggregate interface in a cluster acquires the virtual MAC address that would have been acquired by the first interface in the aggregate.

Link aggregation, HA failover performance, and HA mode

To operate an active-active or active-passive cluster with aggregated interfaces and for best performance of a cluster with aggregated interfaces, the switches used to connect the cluster unit aggregated interfaces together should support configuring multiple Link Aggregation (LAG) groups.

For example, the cluster shown above should be configured into two LAG groups on the external switch: one for the port1 and port2 aggregated interface of FGT_ha_1 and a second one for the port1 and port2 aggregate interface of FGT_ha_2. You should also be able to do the same on the internal switch for the port3 and port4 aggregated interfaces of each cluster unit.

As a result, the subordinate unit aggregated interfaces would participate in LACP negotiation while the cluster is operating. In an active-active mode cluster, packets could be redirected to the subordinate unit interfaces. As well, in active-active or active-passive mode, after a failover the subordinate unit can become a primary unit without having to perform LACP negotiation before it can process traffic. Performing LACP negotiation causes a minor failover delay.

However if you cannot configure multiple LAG groups on the switches, due to the primary and subordinate unit interfaces having the same MAC address, the switch will put all of the interfaces into the same LAG group which would disrupt the functioning of the cluster. To prevent this from happening, you must change the FortiGate aggregated interface configuration to prevent subordinate units from participating in LACP negotiation.

For example, use the following command to prevent subordinate units from participating in LACP negotiation with an aggregate interface named Port1_Port2:

```
config system interface
  edit Port1_Port2
    set lacp-ha-slave disable
  end
```

As a result of this setting, subordinate unit aggregated interfaces cannot accept packets. This means that you cannot operate the cluster in active-active mode because in active-active mode the subordinate units must be able to receive and process packets. Also, failover may take longer because after a failover the subordinate unit has to perform LACP negotiation before being able to process network traffic.

Also, it may also be necessary to configure the switch to use Passive or even Static mode for LACP to prevent the switch from sending packets to the subordinate unit interfaces, which won't be able to process them.

Finally, in some cases depending on the LACP configuration of the switches, you may experience delayed failover if the FortiGate LACP configuration is not compatible with the switch LACP configuration. For example, in some cases setting the FortiGate LACP mode to static reduces the failover delay because the FortiGate does not perform LACP negotiation. However there is a potential problem with this configuration because static LACP does not send periodic LAC Protocol Data Unit (LACPDU) packets to test the connections. So a non-physical failure (for example, if a device is not responding because its too busy) may not be detected and packets could be lost or delayed.

General configuration steps

The section includes GUI and CLI procedures. These procedures assume that the FortiGates are running the same FortiOS firmware build and are set to the factory default configuration.

General configuration steps

1. Apply licenses to the FortiGates to become the cluster.
2. Configure the FortiGates for HA operation.
 - Change each unit's host name.
 - Configure HA.
2. Connect the cluster to the network.
3. View cluster status.
4. Add basic configuration settings and configure the aggregated interfaces.
 - Add a password for the admin administrative account.
 - Add the aggregated interfaces.
 - Disable `lacp-ha-slave` so that the subordinate unit does not send LACP packets.
 - Add a default route.

You could also configure aggregated interfaces in each FortiGate before the units form a cluster.


5. Configure HA port monitoring for the aggregated interfaces.

Configuring active-passive HA cluster that includes aggregated interfaces - GUI

These procedures assume you are starting with two FortiGates with factory default settings.

To configure the FortiGates for HA operation

1. Register and apply licenses to the FortiGate.



2. On the **System Information** dashboard widget, beside **Host Name** select **Change**.
3. Enter a new Host Name for this FortiGate.

New Name	FGT_ha_1
-----------------	----------

4. Select **OK**.
5. Go to **System > HA** and change the following settings.

Mode	Active-Passive	
Group Name	example5.com	
Password	HA_pass_5	
Heartbeat Interface		
	Enable	Priority
port5	Select	50
port6	Select	50

Since port3 and port4 will be used for an aggregated interface, you must change the HA heartbeat configuration to not use those interfaces.

6. Select **OK**.

The FortiGate negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11

- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

7. Power off the first FortiGate.
8. Repeat these steps for the second FortiGate.

Set the second FortiGate host name to:

New Name	FGT_ha_2

To connect the cluster to the network

1. Connect the port1 and port2 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internet. Configure the switch so that the port1 and port2 of FGT_ha_1 make up an aggregated interface and port1 and port2 of FGT_ha_2 make up a second aggregated interface.
2. Connect the port3 and port4 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internal network. Configure the switch so that the port3 and port4 of FGT_ha_1 make up an aggregated interface and port3 and port4 of FGT_ha_2 make up another aggregated interface.
3. Connect the port5 interfaces of FGT_ha_1 and FGT_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port5 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
5. Power on the cluster units.

The units negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete, the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.

1. View the system dashboard.

The HA Status dashboard widget displays how long the cluster has been operating (Uptime) and the time since the last failover occurred (State Changed) You can hover over the State Changed time to see the event that caused the state change You can also click on the HA Status dashboard widget to configure HA settings or to get a listing of the most recent HA events recorded by the cluster.

2. Go to **System > HA** to view the cluster members list.

The list shows both cluster units, their host names, their roles in the cluster, and their device priorities. You can use this list to confirm that the cluster is operating normally. For example, if the list shows only one cluster unit then the other unit has left the cluster for some reason.

To troubleshoot the cluster configuration

See [Troubleshooting HA clusters on page 1458](#) to troubleshoot the cluster.

To add basic configuration settings and the aggregate interfaces

Use the following steps to add a few basic configuration settings.

1. Log into the cluster GUI.
2. Go to **System > Administrators**.
3. Edit **admin** and select **Change Password**.
4. Enter and confirm a new password.
5. Select **OK**.
6. Go to **Network > Static Routes** and temporarily delete the default route.
You cannot add an interface to a aggregated interface if any settings (such as the default route) are configured for it.
7. Go to **Network > Interfaces** and select **Create New > Interface** to add the aggregate interface to connect to the internet.
8. Set **Type** to **802.3ad Aggregate** and configure the aggregate interface to be connected to the internet:

Name	Port1_Port2
Interface Members	port1, port2
IP/Network Mask	172.20.120.141/24

9. Select **OK**.
10. Select **Create New > Interface** to add the aggregate interface to connect to the internal network.
11. Set **Type** to **802.3ad Aggregate** and configure the aggregate interface to be connected to the internet:

Name	Port3_Port4
Interface Members	port3, port4
IP/Netmask	10.11.101.100/24
Administrative Access	HTTPS, PING, SSH

12. Select **OK**.

The virtual MAC addresses of the FortiGate interfaces change to the following. Note that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

13. Connect to the CLI and enter the following command to disable sending LACP packets from the subordinate unit:

```
config system interface
  edit Port1_Port2
    set lacp-ha-slave disable
  next
  edit Port3_Port4
    set lacp-ha-slave disable
end
```

14. Go to **Network > Static Routes**.

15. Add the default route.

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.2
Device	Port1_Port2
Distance	10

16. Select **OK**.

To configure HA port monitoring for the aggregate interfaces

1. Go to **System > HA**.
2. In the cluster members list, edit the primary unit.
3. Configure the following port monitoring for the aggregate interfaces:

Port Monitor	
Port1_Port2	Select
Port3_Port4	Select

4. Select **OK**.

Configuring active-passive HA cluster that includes aggregate interfaces - CLI

These procedures assume you are starting with two FortiGates with factory default settings.

To configure the FortiGates for HA operation

1. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS). All of the FortiGates in a cluster must have the same level of licensing.
2. Install any third-party certificates on the FortiGate.
3. Change the host name for this FortiGate:

```
config system global
    set hostname FGT_ha_1
end
```

4. Configure HA settings.

```
config system ha
    set mode a-p
    set group-name example5.com
    set password HA_pass_5
    set hbdev port5 50 port6 50
end
```

Since port3 and port4 will be used for an aggregated interface, you must change the HA heartbeat configuration.

The FortiGate negotiates to establish an HA cluster. You may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06

- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

4. Repeat these steps for the other FortiGate.

Set the other FortiGate host name to:

```
config system global
    set hostname FGT_ha_2
end
```

To connect the cluster to the network

1. Connect the port1 and port2 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internet. Configure the switch so that the port1 and port2 of FGT_ha_1 make up an aggregated interface and port1 and port2 of FGT_ha_2 make up another aggregated interface.
2. Connect the port3 and port4 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internal network. Configure the switch so that the port3 and port4 of FGT_ha_1 make up an interfaced and port3 and port4 of FGT_ha_2 make up another aggregated interface.

3. Connect the port5 interfaces of FGT_ha_1 and FGT_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port5 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
5. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view cluster status from the CLI.

1. Log into the CLI.
2. Enter `get system status` to verify the HA status of the cluster unit that you logged into. Look for the following information in the command output.

Current HA mode: a-a, master	The cluster units are operating as a cluster and you have connected to the primary unit.
Current HA mode: a-a, backup	The cluster units are operating as a cluster and you have connected to a subordinate unit.
Current HA mode: standalone	The cluster unit is not operating in HA mode

3. Enter the following command to view the status of the cluster:

```
get system ha status
HA Health Status: OK
Model: FortiGate-XXXX
Mode: HA A-P
Group: 0
Debug: 0
Cluster Uptime: 7 days 00:30:26
.
.
.
```

You can use this command to confirm that the cluster is healthy and operating normally, some information about the cluster configuration, and information about how long the cluster has been operating. Information not shown in this example includes how the primary unit was selected, configuration synchronization status, usage stats for each cluster unit, heartbeat status, and the relative priorities of the cluster units.

To troubleshoot the cluster configuration

See [Troubleshooting HA clusters on page 1458](#) to troubleshoot the cluster.

To add basic configuration settings and the aggregate interfaces

Use the following steps to add a few basic configuration settings and the aggregate interfaces.

1. Add a password for the admin administrative account.

```

config system admin
  edit admin
    set password <psswr>
  end

```

2. Temporarily delete the default route.

You cannot add an interface to an aggregate interface if any settings (such as the default route) are configured for it. In this example the index of the default route is 1.

```

config router static
  delete 1
end

```

3. Add the aggregate interfaces:

```

config system interface
  edit Port1_Port2
    set type aggregate
    set lacp-ha-slave disable
    set member port1 port2
    set ip 172.20.120.141/24
    set vdom root
  next
  edit Port3_Port4
    set type aggregate
    set lacp-ha-slave disable
    set member port3 port4
    set ip 10.11.101.100/24
    set vdom root
  end
end

```

The virtual MAC addresses of the FortiGate interfaces change to the following. Note that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10

- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

4. Add the default route.

```
config router static
  edit 1
    set dst 0.0.0.0 0.0.0.0
    set gateway 172.20.120.2
    set device Port1_Port2
  end
```

To configure HA port monitoring for the aggregate interfaces

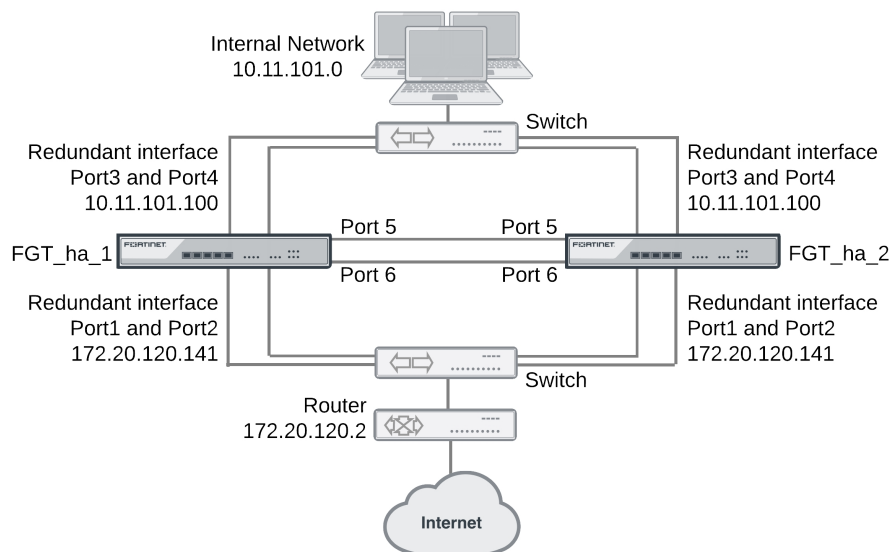
1. Configure HA port monitoring for the aggregate interfaces.

```
config system ha
  set monitor Port1_Port2 Port3_Port4
end
```

Example HA and redundant interfaces

On FortiGate models that support it you can combine two or more interfaces into a single redundant interface. A redundant interface consists of two or more physical interfaces. Traffic is processed by the first physical interface in the redundant interface. If that physical interface fails, traffic fails over to the next physical interface. Redundant interfaces don't have the benefit of improved performance that aggregate interfaces can have, but they do provide failover if a physical interface fails or is disconnected.

Example cluster with a redundant interfaces



This example describes how to configure an HA cluster consisting of two FortiGates with a redundant interface connection to the internet and to an internal network. The connection to the internet uses port1 and port2. The connection to the internal network uses port3 and port4. The HA heartbeat uses port5 and port6.

The redundant interfaces are also configured as HA monitored interfaces.

HA interface monitoring, link failover, and redundant interfaces

HA interface monitoring monitors the redundant interface as a single interface and does not monitor the individual physical interfaces in the redundant interface. HA interface monitoring registers the redundant interface to have failed only if all the physical interfaces in the redundant interface have failed. If only some of the physical interfaces in the redundant interface fail or become disconnected, HA considers the redundant interface to be operating normally.

HA MAC addresses and redundant interfaces

For a standalone FortiGate a redundant interface has the MAC address of the first physical interface added to the redundant interface configuration. A redundant interface consisting of port1 and port2 would have the MAC address of port1.

In an HA cluster, HA changes the MAC addresses of the cluster interfaces to virtual MAC addresses. A redundant interface in a cluster acquires the virtual MAC address that would have been acquired by the first physical interface added to the redundant interface configuration.

Connecting multiple redundant interfaces to one switch while operating in active-passive HA mode

HA assigns the same virtual MAC addresses to the subordinate unit interfaces as are assigned to the corresponding primary unit interfaces. Consider a cluster of two FortiGates operating in active-passive mode with a redundant interface consisting of port1 and port2. You can connect multiple redundant interfaces to the same switch if you configure the switch so that it defines multiple separate redundant interfaces and puts the redundant interfaces of each cluster unit into separate redundant interfaces. In this configuration, each cluster unit forms a separate redundant interface with the switch.

However, if the switch is configured with a single four-port redundant interface configuration, because the same MAC addresses are being used by both cluster units, the switch adds all four interfaces (port1 and port2 from the primary unit and port1 and port2 from the subordinate unit) to the same redundant interface.

To avoid unpredictable results, when you connect a switch to multiple redundant interfaces in an active-passive cluster you should configure separate redundant interfaces on the switch; one for each cluster unit.

Connecting multiple redundant interfaces to one switch while operating in active-active HA mode

In an active-active cluster, all cluster units send and receive packets. To operate a cluster with redundant interfaces in active-active mode, with multiple redundant interfaces connected to the same switch, you must separate the redundant interfaces of each cluster unit into different redundant interfaces on the connecting switch.

General configuration steps

The section includes GUI and CLI procedures. These procedures assume that the FortiGates are running the same FortiOS firmware build and are set to the factory default configuration.

General configuration steps

1. Apply licenses to the FortiGates to become the cluster.
2. Configure the FortiGates for HA operation.

- Change each unit's host name.
 - Configure HA.
2. Connect the cluster to the network.
 3. View cluster status.
 4. Add basic configuration settings and configure the redundant interfaces.
 - Add a password for the admin administrative account.
 - Add the redundant interfaces.
 - Add a default route.

You could also configure redundant interfaces in each FortiGate before they form a cluster.

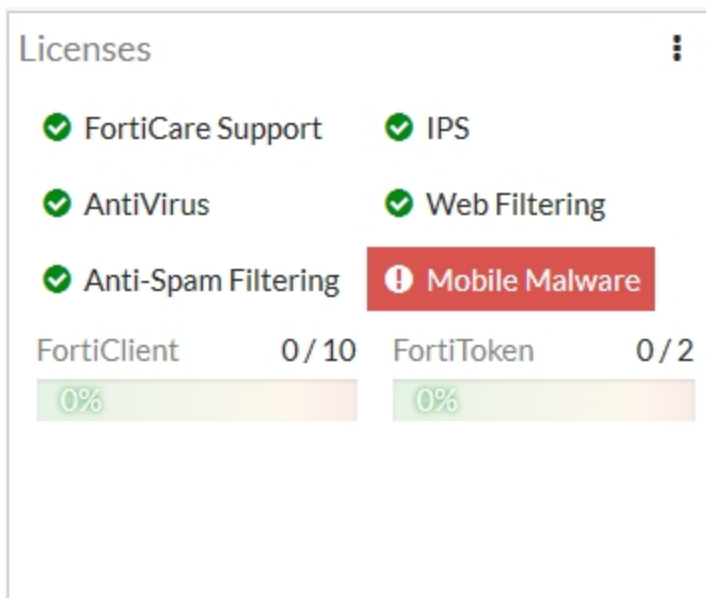
5. Configure HA port monitoring for the redundant interfaces.

Configuring active-passive HA cluster that includes redundant interfaces - GUI

These procedures assume you are starting with two FortiGates with factory default settings.

To configure the FortiGates for HA operation

1. Register and apply licenses to the FortiGate.



2. On the **System Information** dashboard widget, beside **Host Name** select **Change**.
3. Enter a new Host Name for this FortiGate.

New Name	FGT_ha_1
-----------------	----------

4. Select **OK**.
5. Go to **System > HA** and change the following settings.

Mode	Active-Passive
-------------	----------------

Group Name	example6.com	
Password	HA_pass_6	
Heartbeat Interface		
	Enable	Priority
port5	Select	50
port6	Select	50

Since port3 and port4 will be used for a redundant interface, you must change the HA heartbeat configuration.

6. Select OK.

The FortiGate negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following

command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

7. Power off the first FortiGate.
8. Repeat these steps for the second FortiGate.
Set the second FortiGate host name to:

New Name	FGT_ha_2
----------	----------

To connect the cluster to the network

1. Connect the port1 and port2 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internet. Configure the switch so that the port1 and port2 of FGT_ha_1 make up a redundant interface and port1 and port2 of FGT_ha_2 make up another redundant interface.
2. Connect the port3 and port4 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internal network. Configure the switch so that the port3 and port4 of FGT_ha_1 make up a redundant interface and port3 and port4 of FGT_ha_2 make up another redundant interface.
3. Connect the port5 interfaces of FGT_ha_1 and FGT_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port5 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
5. Power on the cluster units.
The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.

1. View the system dashboard.
The HA Status dashboard widget displays how long the cluster has been operating (Uptime) and the time since the last failover occurred (State Changed). You can hover over the State Changed time to see the event that caused the state change. You can also click on the HA Status dashboard widget to configure HA settings or to get a listing of the most recent HA events recorded by the cluster.
2. Go to **System > HA** to view the cluster members list.
The list shows both cluster units, their host names, their roles in the cluster, and their device priorities.

You can use this list to confirm that the cluster is operating normally. For example, if the list shows only one cluster unit then the other unit has left the cluster for some reason.

To troubleshoot the cluster configuration

See [Troubleshooting HA clusters on page 1458](#).

To add basic configuration settings and the redundant interfaces

Use the following steps to add a few basic configuration settings.

1. Log into the cluster GUI.
2. Go to **System > Administrators**.
3. Edit **admin** and select **Change Password**.
4. Enter and confirm a new password.
5. Select **OK**.
6. Go to **Network > Static Routes** and temporarily delete the default route.
You cannot add an interface to a redundant interface if any settings (such as the default route) are configured for it.
7. Go to **Network > Interfaces** and select **Create New > Interface** to add the redundant interface to connect to the internet.
8. Set **Type** to **Redundant Interface** and configure the redundant interface to be connected to the internet:

Name	Port1_Port2
Physical Interface Members	port1, port2
IP/Netmask	172.20.120.141/24

9. Select **OK**.
10. Select **Create New** to add the redundant interface to connect to the internal network.
11. Set **Type** to **Redundant Interface** and configure the redundant interface to be connected to the internet:

Name	Port3_Port4
Physical Interface Members	port3, port4
IP/Netmask	10.11.101.100/24
Administrative Access	HTTPS, PING, SSH

12. Select **OK**.
The virtual MAC addresses of the FortiGate interfaces change to the following. Note that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:
 - port1 interface virtual MAC: 00-09-0f-09-00-00
 - port10 interface virtual MAC: 00-09-0f-09-00-01
 - port11 interface virtual MAC: 00-09-0f-09-00-02

- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

13. Go to **Router > Static > Static Routes**.

14. Add the default route.

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.2
Device	Port1_Port2
Distance	10

15. Select **OK**.

To configure HA port monitoring for the redundant interfaces

1. Go to **System > HA**.
2. In the cluster members list, edit the primary unit.
3. Configure the following port monitoring for the redundant interfaces:

	Port Monitor
Port1_Port2	Select
Port3_Port4	Select

4. Select **OK**.

Configuring active-passive HA cluster that includes redundant interfaces - CLI

These procedures assume you are starting with two FortiGates with factory default settings.

To configure the FortiGates for HA operation

1. Register and apply licenses to the FortiGate. This includes **FortiCloud** activation and **FortiClient** licensing, and entering a license key if you purchased more than 10 **Virtual Domains** (VDOMS). All of the FortiGates in a cluster must have the same level of licensing.
2. You can also install any third-party certificates on the primary FortiGate before forming the cluster. Once the cluster is formed third-party certificates are synchronized to the backup FortiGate.
We recommend that you add FortiToken licenses and FortiTokens to the primary unit after the cluster has formed.
3. Change the host name for this FortiGate:

```
config system global
    set hostname FGT_ha_1
end
```

4. Configure HA settings.

```
config system ha
    set mode a-p
    set group-name example6.com
    set password HA_pass_6
    set hbdev port5 50 port6 50
end
```

Since port3 and port4 will be used for a redundant interface, you must change the HA heartbeat configuration.

The FortiGate negotiates to establish an HA cluster. You may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

4. Repeat these steps for the other FortiGate.

Set the other FortiGate host name to:

```
config system global
    set hostname FGT_ha_2
end
```

To connect the cluster to the network

1. Connect the port1 and port2 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internet. Configure the switch so that the port1 and port2 of FGT_ha_1 make up a redundant interface and port1 and port2 of FGT_ha_2 make up another redundant interface.
2. Connect the port3 and port4 interfaces of FGT_ha_1 and FGT_ha_2 to a switch connected to the internal network. Configure the switch so that the port3 and port4 of FGT_ha_1 make up a redundant interface and port3 and port4 of FGT_ha_2 make up another redundant interface.
3. Connect the port5 interfaces of FGT_ha_1 and FGT_ha_2 together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
4. Connect the port5 interfaces of the cluster units together. You can use a crossover Ethernet cable or regular Ethernet cables and a switch.
5. Power on the cluster units.
The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention and normally takes less than a minute.

When negotiation is complete the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view cluster status from the CLI.

1. Log into the CLI.
2. Enter `get system status` to verify the HA status of the cluster unit that you logged into. Look for the following information in the command output.

Current HA mode: a-a, master	The cluster units are operating as a cluster and you have connected to the primary unit.
Current HA mode: a-a, backup	The cluster units are operating as a cluster and you have connected to a subordinate unit.
Current HA mode: standalone	The cluster unit is not operating in HA mode

3. Enter the following command to confirm the HA configuration of the cluster:

```
get system ha status
HA Health Status: OK
Model: FortiGate-XXXX
Mode: HA A-P
Group: 0
Debug: 0
Cluster Uptime: 7 days 00:30:26
.
.
.
```

You can use this command to confirm that the cluster is healthy and operating normally, some information about the cluster configuration, and information about how long the cluster has been operating. Information not shown in this example includes how the primary unit was selected, configuration synchronization status, usage stats for each cluster unit, heartbeat status, and the relative priorities of the cluster units.

To troubleshoot the cluster configuration

See [Troubleshooting HA clusters on page 1458](#).

To add basic configuration settings and the redundant interfaces

Use the following steps to add a few basic configuration settings and the redundant interfaces.

1. Add a password for the admin administrative account.

```
config system admin
  edit admin
    set password <psswr>
  end
```

2. Temporarily delete the default route.

You cannot add an interface to a redundant interface if any settings (such as the default route) are configured for it. In this example the index of the default route is 1.

```
config router static
  delete 1
end
```

3. Add the redundant interfaces:

```
config system interface
  edit Port1_Port2
    set type redundant
    set member port1 port2
```

```

        set ip 172.20.120.141/24
        set vdom root
    next
    edit Port3_Port4
        set type redundant
        set member port3 port4
        set ip 10.11.101.100/24
        set vdom root
    end

```

The virtual MAC addresses of the FortiGate interfaces change to the following. Note that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

4. Add the default route.

```

config router static
    edit 1
        set dst 0.0.0.0 0.0.0.0
        set gateway 172.20.120.2
        set device Port1_Port2
    end

```

To configure HA port monitoring for the redundant interfaces

1. Configure HA port monitoring for the redundant interfaces.

```

config system ha
    set monitor Port1_Port2 Port3_Port4
end

```


Troubleshooting HA clusters

This section describes some HA clustering troubleshooting techniques.

Ignoring hardware revisions

Many FortiGate platforms have gone through multiple hardware versions and in some cases the hardware changes prevent cluster formation. If you run into this problem you can use the following command on each FortiGate to cause the cluster to ignore different hardware versions:

```
execute ha ignore-hardware-revision enable
```

This command is only available on FortiGates that have had multiple hardware revisions.

By default the command is set to prevent cluster formation between FortiGates with different hardware revisions. You can enter the following command to view its status:

```
execute ha ignore-hardware-revision status
```

Usually the incompatibility is caused by different hardware versions having different hard disks and enabling this command disables the hard disks in each FortiGate. As a result of disabling hard disks the cluster will not support logging to the hard disk or WAN Optimization.

If the FortiGates do have compatible hardware versions or if you want to run a FortiGate in standalone mode you can enter the following command to disable ignoring the hardware revision and enable the hard disks:

```
execute ha ignore-hardware-revision disable
```

Affected models include but are not limited to:

- FortiGate-100D
- FortiGate-300C
- FortiGate-600C
- FortiGate-800C
- FortiGate-80C and FortiWiFi-80C
- FortiGate-60C



It's possible that a cluster will not form because the disk partition sizes of the cluster units are different. You can use the `diagnose sys ha checksum test | grep storage` command to check the disk storage checksum of each cluster unit. If the checksums are different then visit the [Fortinet Support](#) website for help in setting up compatible storage partitions.

Before you set up a cluster

Before you set up a cluster ask yourself the following questions about the FortiGates that you are planning to use to create a cluster.

1. Do all the FortiGates have the same hardware configuration? Including the same hard disk configuration?
2. Do all of the FortiGates have the same FortiGuard, FortiCloud, FortiClient, VDOM and FortiOS Carrier licensing?
3. Do all the FortiGates have the same firmware build?
4. Are all the FortiGates set to the same operating mode (NAT or transparent)?
5. Are all the FortiGates operating in single VDOM mode?
6. If the FortiGates are operating in multiple VDOM mode do they all have the same VDOM configuration?



In some cases you may be able to form a cluster if different FortiGates have different firmware builds, different VDOM configurations, and are in different operating modes. However, if you encounter problems they may be resolved by installing the same firmware build on each unit, and give them the same VDOM configuration and operating mode. If the FortiGates in the cluster have different licenses, the cluster will form but it will operate with the lowest licensing level.

Troubleshooting the initial cluster configuration

This section describes how to check a cluster when it first starts up to make sure that it is configured and operating correctly. This section assumes you have already configured your HA cluster.

To verify that a cluster can process traffic and react to a failure

1. Add a basic security policy configuration and send network traffic through the cluster to confirm connectivity.
For example, if the cluster is installed between the internet and an internal network, set up a basic internal to external security policy that accepts all traffic. Then from a PC on the internal network, browse to a website on the internet or ping a server on the internet to confirm connectivity.
2. From your management PC, set ping to continuously ping the cluster, and then start a large download, or in some other way establish ongoing traffic through the cluster.
3. While traffic is going through the cluster, disconnect the power from one of the cluster units.
You could also shut down or restart a cluster unit.
Traffic should continue with minimal interruption.
4. Start up the cluster unit that you disconnected.
The unit should re-join the cluster with little or no affect on traffic.
5. Disconnect a cable from one of the HA heartbeat interfaces.
The cluster should keep functioning, using the other HA heartbeat interface.
6. If you have port monitoring enabled, disconnect a network cable from a monitored interface.
Traffic should continue with minimal interruption.

To verify the cluster configuration from the GUI

Use these steps if a cluster is formed just to verify its status and configuration.

1. Log into the cluster GUI.
2. Check the system dashboard to verify that the System Information widget displays all of the cluster units.
3. Check the Unit Operation widget graphic to verify that the correct cluster unit interfaces are connected.
4. Go to **System > HA** or from the System Information dashboard widget select **HA Status > Configure** and verify that all of the cluster units are displayed on the HA Cluster list.
5. From the cluster members list, edit the primary unit (master) and verify the cluster configuration is as expected.

To troubleshoot the cluster configuration from the GUI

Use these steps if the FortiGates don't successfully form a cluster:

1. Connect to each cluster unit GUI and verify that the HA configurations are the same. The HA configurations of all of the cluster units must be identical. Even though the HA configuration is very simple you can easily make a small

mistake that prevents a FortiGate from joining a cluster.

2. If the configurations are the same, try re-entering the HA **Password** on each cluster unit in case you made an error typing the password when configuring one of the cluster units.
3. Check that the correct interfaces of each cluster unit are connected.
Check the cables and interface LEDs.

Use the Unit Operation dashboard widget, system network interface list, or cluster members list to verify that each interface that should be connected actually is connected.

If the link is down re-verify the physical connection. Try replacing network cables or switches as required.

To verify the cluster configuration from the CLI

Use these steps if a cluster is formed just to verify its status and configuration.

1. Log into each cluster unit CLI.
You can use the console connection if you need to avoid the problem of units having the same IP address.
2. Enter the command `get system status`.
Look for the following information in the command output.

Current HA mode: a-a, master	The cluster units are operating as a cluster and you have connected to the primary unit.
Current HA mode: a-a, backup	The cluster units are operating as a cluster and you have connected to a subordinate unit.
Current HA mode: standalone	The cluster unit is not operating in HA mode

3. Verify that the `get system ha status` command shows that the cluster health is OK and shows that all of the cluster units have joined the cluster.
4. Enter the `get system ha` command to verify that the HA configuration is correct and the same for each cluster unit.

To troubleshoot the cluster configuration from the CLI

Try these steps if the FortiGates don't successfully form a cluster:

1. Try using the following command to re-enter the cluster password on each cluster unit in case you made an error typing the password when configuring one of the cluster units.

```
config system ha
  set password <password>
end
```

2. Check that the correct interfaces of each cluster unit are connected.
Check the cables and interface LEDs.

Use `get hardware nic <interface_name>` command to confirm that each interface is connected. If the interface is connected the command output should contain a `Link: up` entry similar to the following:

```

get hardware nic port1
.
.
.
Link: up
.
.
.

```

If the link is down, re-verify the physical connection. Try replacing network cables or switches as required.

More troubleshooting information

Much of the information in this HA guide can be useful for troubleshooting HA clusters. Here are some links to sections with more information.

- If sessions are lost after a failover you may need to change route-ttl to keep synchronized routes active longer. See [Synchronizing kernel routing tables on page 1546](#)
- To control which cluster unit becomes the primary unit, you can change the device priority and enable override. See [Controlling primary unit selection using device priority and override on page 1395](#)
- Changes made to a cluster can be lost if override is enabled. See [Configuration changes can be lost if override is enabled on page 1396](#)
- When override is enabled, after a failover traffic may be disrupted if the primary unit rejoins the cluster before the session tables are synchronized or for other reasons such as if the primary unit is configured for DHCP or PPPoE. See [Delaying how quickly the primary unit rejoins the cluster when override is enabled on page 1397](#).
- In some cases, age differences among cluster units result in the wrong cluster unit becoming the primary unit. For example, if a cluster unit set to a high priority reboots, that unit will have a lower age than other cluster units. You can resolve this problem by resetting the age of one or more cluster units. See [Primary unit selection and age on page 1387](#) You can also adjust how sensitive the cluster is to age differences. This can be useful if large age differences cause problems. See [Cluster age difference margin \(grace period\) on page 1388](#) and [Changing the cluster age difference margin on page 1388](#).
- If one of the cluster units needs to be serviced or removed from the cluster for other reasons, you can do so without affecting the operation of the cluster. See [Disconnecting a cluster unit from a cluster on page 1518](#)
- The GUI and CLI will not allow you to configure HA if you have enabled FGSP HA. See [FortiGate Session Life Support Protocol \(FGSP\) on page 1606](#).
- The GUI and CLI will not allow you to configure HA if one or more FortiGate interfaces is configured as a PPTP or L2TP client.
- The FGCP is compatible with DHCP and PPPoE but care should be taken when configuring a cluster that includes a FortiGate interface configured to get its IP address with DHCP or PPPoE. Fortinet recommends that you turn on DHCP or PPPoE addressing for an interface after the cluster has been configured. See [FortiGate HA compatibility with DHCP and PPPoE on page 1397](#).
- Some third-party network equipment may prevent HA heartbeat communication, resulting in a failure of the cluster or the creation of a split brain scenario. For example, some switches use packets with the same Ethertype as HA heartbeat packets use for internal functions and when used for HA heartbeat communication the switch generates CRC errors and the packets are not forwarded. See [Heartbeat packet Ethernets on page 1527](#).
- Very busy clusters may not be able to send HA heartbeat packets quickly enough, also resulting in a split brain scenario. You may be able to resolve this problem by modifying HA heartbeat timing. See [Modifying heartbeat timing on page 1528](#).

- Very busy clusters may suffer performance reductions if session pickup is enabled. If possible you can disable this feature to improve performance. If you require session pickup for your cluster, several options are available for improving session pickup performance. See [Improving session synchronization performance on page 1](#).
- If it takes longer than expected for a cluster to failover you can try changing how the primary unit sends gratuitous ARP packets. See [Changing how the primary unit sends gratuitous ARP packets after a failover on page 1531](#).
- You can also improve failover times by configuring the cluster for sub-second failover. See [Sub-second failover on page 1555](#) and [Failover performance on page 1566](#).
- When you first put a FortiGate in HA mode you may lose connectivity to the unit. This occurs because HA changes the MAC addresses of all FortiGate interfaces, including the one that you are connecting to. The cluster MAC addresses also change if you change some HA settings such as the cluster group ID. The connection will be restored in a short time as your network and PC updates to the new MAC address. To reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.
- Since HA changes all cluster unit MAC addresses, if your network uses MAC address filtering you may have to make configuration changes to account for the HA MAC addresses.
- A network may experience packet loss when two FortiGate HA clusters have been deployed in the same broadcast domain. Deploying two HA clusters in the same broadcast domain can result in packet loss because of MAC address conflicts. The packet loss can be diagnosed by pinging from one cluster to the other or by pinging both of the clusters from a device within the broadcast domain. You can resolve the MAC address conflict by changing the HA Group ID configuration of the two clusters. The HA Group ID is sometimes also called the Cluster ID. See [Diagnosing packet loss with two FortiGate HA clusters in the same broadcast domain on page 1535](#).
- The cluster CLI displays `slave is not in sync` messages if there is a synchronization problem between the primary unit and one or more subordinate units. See [How to diagnose HA out of sync messages on page 1543](#).
- If you have configured dynamic routing and the new primary unit takes too long to update its routing table after a failover you can configure graceful restart and also optimize how routing updates are synchronized. See [Configuring graceful restart for dynamic routing failover on page 1547](#) and [Synchronizing kernel routing tables on page 1546](#).
- Some switches may not be able to detect that the primary unit has become a subordinate unit and will keep sending packets to the former primary unit. This can occur after a link failover if the switch does not detect the failure and does not clear its MAC forwarding table. See [Updating MAC forwarding tables when a link failover occurs on page 1553](#).
- If a link not directly connected to a cluster unit (for example, between a switch connected to a cluster interface and the network) fails you can enable remote link failover to maintain communication. See [Remote link failover on page 1556](#).
- If you find that some cluster units are not running the same firmware build you can reinstall the correct firmware build on the cluster to upgrade all cluster units to the same firmware build. See [Synchronizing the firmware build running on a new cluster unit on page 1508](#).

Virtual clusters

This chapter describes virtual clustering, a variation of the FGCP for FortiGates with multiple VDOMs. Virtual clustering includes can operate in active-passive or active-active HA mode for clusters of up to four FortiGates. Active-passive virtual clustering includes VDOM partitioning that to distribute traffic for different VDOMs between the primary and backup FortiGates.

Virtual clustering overview

Virtual clustering is an extension of the FGCP that provides failover protection between two instances of one or more VDOMs operating on two FortiGates in a virtual cluster.

A standard virtual cluster consists of up to four FortiGates operating in active-passive or active-active HA mode with multiple VDOMS enabled.

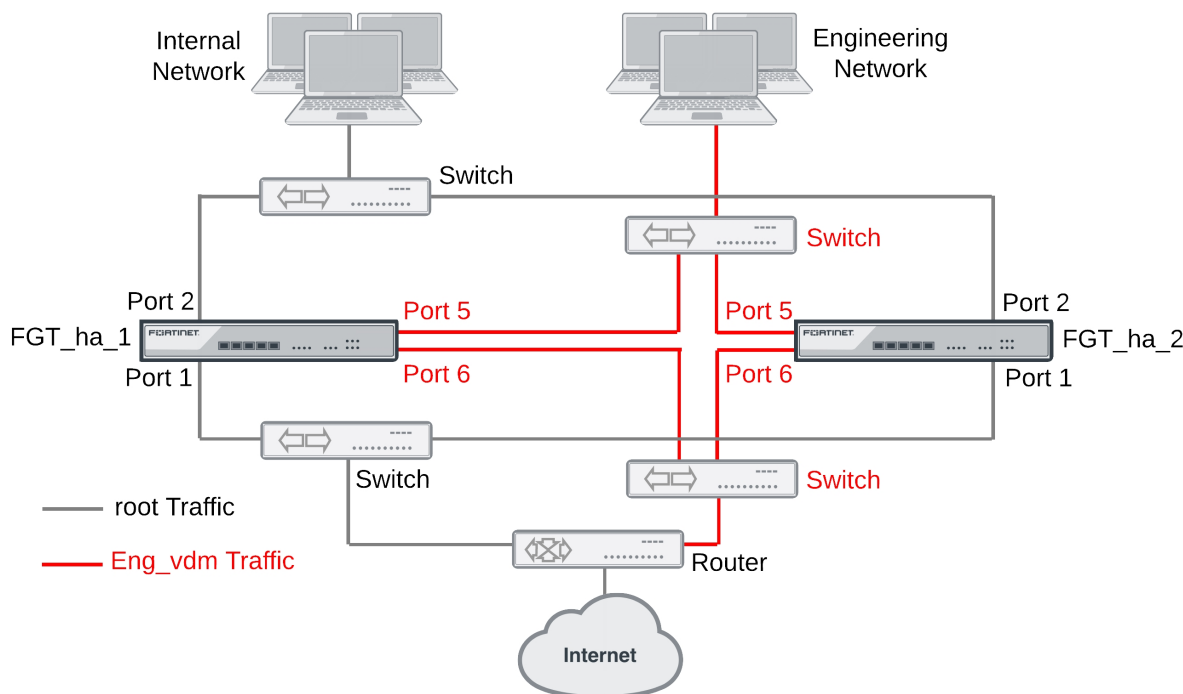
Active-passive virtual clustering uses VDOM partitioning to send traffic for some VDOMs to the primary FortiGate and traffic for other VDOMs to the backup FortiGate. Traffic distribution between both FortiGates can potentially improve throughput. If a failure occurs and only one FortiGate continues to operate, all traffic fails over to that FortiGate, similar to normal HA. If the failed FortiGates restart, the configured traffic distribution is restored.

Active-active virtual clustering operates just the same as standard FGCP active-active HA, distributing traffic to all of the FortiGates in the cluster using FGCP load balancing.

In an active-passive virtual cluster of two FortiGates, the primary and backup FortiGates share traffic processing according to the VDOM partitioning configuration. If you add a third or fourth FortiGate, the primary and first backup FortiGate process all traffic and the other one or two FortiGates operate in standby mode. If the primary or first backup FortiGate fails, one of the other FortiGates becomes the new primary or backup FortiGate and begins processing traffic.

The figure below shows an example virtual cluster configuration consisting of two FortiGates. The virtual cluster has two VDOMs, root and Eng_vdm.

Example virtual cluster



The root VDOM includes the port1 and port2 interfaces. The Eng_vdm VDOM includes the port5 and port6 interfaces. The port3 and port4 interfaces (not shown in the diagram) are the HA heartbeat interfaces.



If you don't want active-passive virtual clustering to distribute traffic between FortiGates, you can configure VDOM partitioning to send traffic for all VDOMs to the primary unit. The result is the same as standard active-passive FCGP HA, all traffic is processed by the primary FortiGate.

Separation of VDOM traffic

Virtual clustering creates a cluster between instances of each VDOM on the two FortiGates in the virtual cluster. All traffic to and from a given VDOM is sent to one of the FortiGates where it stays within its VDOM and is only processed by that VDOM. One FortiGate is the primary FortiGate for each VDOM and one FortiGate is the backup FortiGate for each VDOM. The primary FortiGate processes all traffic for its VDOMs. The backup FortiGate processes all traffic for its VDOMs.



If your cluster has a VLAN that is part of a different VDOM than the physical interface that the VLAN has been added to, then you must configure VDOM partitioning to keep traffic for both of these VDOMs on the same FortiGate.

Virtual clustering and heartbeat interfaces

The HA heartbeat provides the same HA services in a virtual clustering configuration as in a standard HA configuration. One set of HA heartbeat interfaces provides HA heartbeat services for all of the VDOMs in the cluster. You do not have to add a heartbeat interface for each VDOM.

Virtual clustering and load balancing

There are two ways to configure load balancing for virtual clustering. The first is to set the HA mode to active-active. The second is to configure VDOM partitioning. For virtual clustering, setting the HA Mode to active-active has the same result as active-active HA for a cluster without virtual domains. The primary FortiGate receives all sessions and load balances them among the cluster units according to the load balancing schedule. All cluster units process traffic for all virtual domains.

In an active-passive virtual clustering configuration, you can configure a form of load balancing by using VDOM partitioning to distribute traffic between the primary and backup FortiGates. While a cluster is operating, you can change the VDOM partitioning configuration to change the distribution of traffic between the cluster units. For example, if you have two VDOMs with high traffic volume you can set up VDOM partitioning so that different FortiGates process the traffic for each high-volume VDOM. If over time traffic patterns change you can dynamically re-adjust VDOM partitioning to optimize traffic throughput. VDOM partitioning can be changed at any time with only minor traffic disruptions.

Configuring virtual clustering

Configuring virtual clustering is the same as configuring standard FCGP HA with the addition of VDOM partitioning. Using VDOM partitioning you can control the distribution of VDOMs, and the traffic they process, between the FortiGates in the cluster.

VDOM partitioning can be thought of in two parts. First there is configuring the distribution of VDOMs between two virtual clusters. By default, all VDOMs are in virtual cluster 1 and virtual cluster 1 is associated with the primary FortiGate. In this configuration, the primary FortiGate processes all traffic. If you want traffic to be processed by the backup FortiGate, you need to enable virtual cluster 2, move some of the VDOMs to it, and associate virtual cluster 2 with the backup FortiGate.



Since there are only two virtual clusters, even in a virtual clustering configuration of three or four FortiGates only two of the FortiGates process traffic. The third and fourth FortiGates operate in standby mode and process traffic after a failover.

By default all VDOMs are in virtual cluster 1 and the primary FortiGate processes all traffic.

You associate a virtual cluster with a FortiGate using priorities. The FortiGate with the highest device priority is associated with virtual cluster 1. To associate a FortiGate with virtual cluster 2 you must enable virtual cluster 2 and set the virtual cluster 2 device priority. The FortiGate with the highest virtual cluster 2 device priority processes traffic for the VDOMs added to virtual cluster 2. (Reminder: device priorities are not synchronized.)



If both FortiGates have the same device priority, virtual cluster 1 is associated with the primary FortiGate. If both FortiGates have the same virtual cluster 2 device priority, virtual cluster 2 is associated with the primary FortiGate.

Virtual clustering and the override setting

Enabling virtual cluster 2 also turns on the HA override setting. Enabling override is required for virtual clustering to function as configured. Enabling override causes the cluster to negotiate every time a failure occurs. If override is not enabled, the cluster will not negotiate after all failures. While more frequent negotiation may cause more minor traffic disruptions, with virtual clustering its more important to negotiate after any failure to make sure the correct traffic flows are maintained.

Example virtual clustering configuration

For example, consider a configuration that includes four VDOMs: root, Engineering, Marketing, and Finance. You can use the following configuration to send root and Engineering traffic to the primary FortiGate and Marketing and Finance traffic to the backup FortiGate.

First, on the primary FortiGate:

- Set the device priority to 200
- Enable virtual cluster 2 (vcluster2)
- Set the virtual cluster 2 device priority (secondary-vcluster) to 50
- Add the Marketing and Finance VDOMs to virtual cluster 2 (secondary-vcluster)



When you enable multiple VDOMs, virtual cluster 2 is enabled by default. Even so the command to enable virtual cluster 2 is included in this example in case for some reason it has been disabled. Enabling virtual cluster 2 also enables override.

```
config global
  config system ha
    set mode a-p
    set group-name mygroup
    set password <password>
    set priority 200
    set vcluster2 enable
    config secondary-vcluster
      set vdom Marketing Finance
      set priority 50
    end
  end
end
```

Then on the backup FortiGate:

- Set the device priority to 50 (lower than the primary FortiGate)
- Enable virtual cluster 2 (vcluster2)
- Set the virtual cluster 2 device priority (secondary-vcluster) to 200 (higher than the primary FortiGate).

```
config global
  config system ha
    set mode a-p
    set group-name mygroup
    set password <password>
    set priority 50
    set vcluster2 enable
    config secondary-vcluster
      set priority 200
    end
  end
end
```



Since the primary FortiGate has the highest device priority, the primary unit processes all traffic for the VDOMs in virtual cluster 1. Since the backup FortiGate has the highest virtual cluster 2 device priority, the backup FortiGate processes all traffic for the VDOMs in virtual cluster 2. The primary FortiGate configuration adds the VDOMs to virtual cluster 2. All you have to configure on the backup FortiGate for virtual cluster 2 is the virtual cluster 2 (or secondary-vcluster) device priority.

Adding a third FortiGate to the virtual cluster

You can add a third FortiGate to the virtual cluster and configure it so that if the primary FortiGate fails, the third FortiGate becomes the new primary FortiGate or if the backup FortiGate fails, the third FortiGate becomes the new backup FortiGate.

On the third FortiGate:

- Set the device priority to 150 (lower than the primary FortiGate but higher than the backup FortiGate)
- Enable virtual cluster 2 (vcluster2)
- Set the virtual cluster 2 device priority (secondary-vcluster) to 100 (higher than the primary FortiGate but lower than the backup FortiGate)

```
config global
  config system ha
    set mode a-p
    set group-name mygroup
    set password <password>
    set priority 150
    set vcluster2 enable
    config secondary-vcluster
      set priority 100
    end
  end
end
```

Adding a fourth FortiGate to the virtual cluster

You can add a fourth FortiGate to the virtual cluster and configure it so that:

- If the primary FortiGate fails, the third FortiGate becomes the new primary FortiGate, the backup FortiGate continues to operate as the backup FortiGate.
- If the backup FortiGate fails, the fourth FortiGate becomes the new backup FortiGate.
- If both the primary and backup FortiGates fail, the third FortiGate becomes the primary FortiGate and the fourth FortiGate becomes the backup FortiGate.

On the fourth FortiGate:

- Set the device priority to 100 (lower than the primary and third FortiGate but higher than the backup FortiGate)
- Enable virtual cluster 2 (vcluster2)
- Set the virtual cluster 2 device priority (secondary-vcluster) to 150 (higher than the primary FortiGate and the third FortiGate but lower than the backup FortiGate)

```
config global
  config system ha
    set mode a-p
    set group-name mygroup
    set password <password>
    set priority 100
```

```
set vcluster2 enable
config secondary-vcluster
    set priority 150
end
end
```

Virtual clustering with four FortiGates recommended configuration

As described in the previous sections, here is a recommended device priority configuration for a virtual cluster consisting of four FortiGates. Other configurations are also supported depending on how you want the virtual cluster to respond to a failure.

FortiGate	Device Priority	Virtual Cluster 2 Device Priority
Primary	200	50
Backup	50	100
Third	150	200
Fourth	100	150

Virtual clustering GUI configuration

From the GUI, you configure virtual clustering from the Global menu by going to **System > HA** setting the **Mode** to **Active-Passive** and enabling **VDOM Partitioning**.

Example primary FortiGate virtual clustering configuration

Mode	Active-Passive ▼
Device priority ⓘ	200

Cluster Settings

Group name	My-vcluster	
Password	●●●●●●●●	Change
Session pickup	<input type="checkbox"/>	
Monitor interfaces	<div> wan1 × </div> <div>+</div>	
Heartbeat interfaces	<div> lan4 × </div> <div> lan5 × </div> <div>+</div>	

Heartbeat Interface Priority ⓘ

lan4	<input type="range"/>	200
lan5	<input type="range"/>	100

☐ Management Interface Reservation☒ VDOM Partitioning

Virtual cluster 1	<div> root × </div> <div> Engineering × </div> <div>+</div>	
Virtual cluster 2	<div> Finance × </div> <div> Marketing × </div> <div>+</div>	

Virtual clustering configuration examples

See the following cookbook recipe for a virtual clustering configuration example. This example shows how to set up a virtual cluster of two FortiGates and then how to add a third and fourth FortiGate to the virtual cluster configuration.

- FGCP virtual clustering with two FortiGates
- FGCP virtual clustering with two FortiGates

Inter-VDOM links in a virtual clustering configuration

In a virtual domain configuration you can use inter-VDOM links to route traffic between two virtual domains operating in a single FortiGate without using physical interfaces. Adding an inter-VDOM link has the affect of adding two interfaces to the FortiGate and routing traffic between the virtual domains using the inter-VDOM link interfaces.

In a virtual clustering configuration inter-VDOM links can only be made between virtual domains that are in the same virtual cluster. So, if you are planning on configuring inter-VDOM links in a virtual clustering configuration, you should make sure the virtual domains that you want to link are in the same virtual cluster.

For example, the following tables show an example virtual clustering configuration where each virtual cluster contains four virtual domains. In this configuration you can configure inter-VDOM links between root and vdom_1 and between vdom_2 and vdom_3. But, you cannot configure inter-VDOM links between root and vdom_2 or between vdom_1 and vdom_3 (and so on).

Hostname		
Virtual Domains	FortiGate_A	FortiGate_B
root	Priority	Priority
	200	100
vdom_1	Role	Role
	Primary	Subordinate

Hostname		
Virtual Domains	FortiGate_A	FortiGate_B
vdom_2	Priority	Priority
	100	200
vdom_3	Role	Role
	Subordinate	Primary

Configuring inter-VDOM links in a virtual clustering configuration

Configuring inter-VDOM links in a virtual clustering configuration is very similar to configuring inter-VDOM links for a standalone FortiGate. The main difference the `config system vdom-link` command includes the `vcluster` keyword. The default setting for `vcluster` is `vcluster1`. So you only have to use the `vcluster` keyword if you are added an inter-VDOM link to virtual cluster 2.

To add an inter-VDOM link to virtual cluster 1

This procedure describes how to create an inter-VDOM link to virtual cluster 1 that results in a link between the root and vdom_1 virtual domains.



Inter-VDOM links are also called internal point-to-point interfaces.

1. Add an inter-VDOM link called `vc1link`.

```
config global
  config system vdom-link
    edit vc1link
  end
```

Adding the inter-VDOM link also adds two interfaces. In this example, these interfaces are called `vc1link0` and `vc1link1`. These interfaces appear in all CLI and GUI interface lists. These interfaces can only be added to virtual domains in virtual cluster 1.

2. Bind the `vc1link0` interface to the root virtual domain and bind the `vc1link1` interface to the `vdom_1` virtual domain.

```
config system interface
  edit vc1link0
    set vdom root
  next
  edit vc1link1
    set vdom vdom_1
  end
```

To add an inter-VDOM link to virtual cluster 2

This procedure describes how to create an inter-VDOM link to virtual cluster 2 that results in a link between the `vdom_2` and `vdom_3` virtual domains.

1. Add an inter-VDOM link called `vc2link`.

```
config global
  config system vdom-link
    edit vc2link
      set vcluster vcluster2
    end
```

Adding the inter-VDOM link also adds two interfaces. In this example, these interfaces are called `vc2link0` and `vc2link1`. These interfaces appear in all CLI and GUI interface lists. These interfaces can only be added to virtual domains in virtual cluster 2.

2. Bind the `vc2link0` interface to the `vdom_2` virtual domain and bind the `vc2link1` interface to the `vdom_3` virtual domain.

```
config system interface
  edit vc2link0
    set vdom vdom_2
  next
  edit vc2link1
    set vdom vdom_3
  end
```

Troubleshooting virtual clustering

Troubleshooting virtual clusters is similar to troubleshooting any cluster (see [FGCP configuration examples and troubleshooting on page 1412](#)). This section describes a few testing and troubleshooting techniques for virtual clustering.

To test the VDOM partitioning configuration

You can do the following to confirm that traffic for different VDOMs will be distributed among both FortiGates in the virtual cluster. These steps assume the cluster is otherwise operating correctly.

1. Log into the GUI or CLI using the IP addresses of interfaces in each VDOM.
Confirm that you have logged into the FortiGate that should be processing traffic for that VDOM by checking the HTML title displayed by your web browser or the CLI prompt. Both of these should include the host name of the cluster unit that you have logged into. Also on the system Dashboard, the System Information widget displays the serial number of the FortiGate that you logged into. From the CLI the `get system status` command displays the status of the cluster unit that you logged into.
2. To verify that the correct cluster unit is processing traffic for a VDOM:
 - Add security policies to the VDOM that allow communication between the interfaces in the VDOM.
 - Optionally enable traffic logging and other monitoring for that VDOM and these security policies.
 - Start communication sessions that pass traffic through the VDOM.
 - Log into the GUI and go to **System > HA**. Verify that the statistics display shows more active sessions, total packets, network utilization, and total bytes for the unit that should be processing all traffic for the VDOM.
 - Optionally check traffic logging and the Top Sessions Widget for the FortiGate that should be processing traffic for that VDOM to verify that the traffic is being processed by this FortiGate.

Full mesh HA

This chapter provides an introduction to full mesh HA and also contains general procedures and configuration examples that describe how to configure FortiGate full mesh HA.

The examples in this chapter include example values only. In most cases you will substitute your own values. The examples in this chapter also do not contain detailed descriptions of configuration parameters.

Full mesh HA overview

When two or more FortiGates are connected to a network in an HA cluster the reliability of the network is improved because the HA cluster replaces a single FortiGate as a single point of failure. With a cluster, a single FortiGate is replaced by a cluster of two or more FortiGates.

However, even with a cluster, potential single points of failure remain. The interfaces of each cluster unit connect to a single switch and that switch provides a single connection to the network. If the switch fails or if the connection between the switch and the network fails service is interrupted to that network.

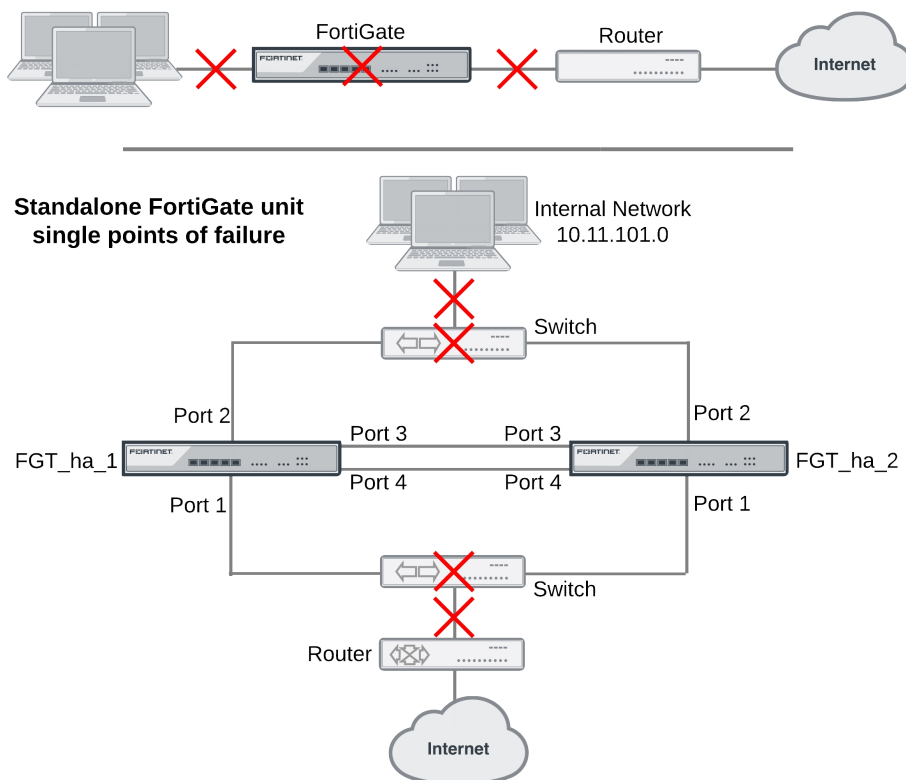
The HA cluster does improve the reliability of the network because switches are not as complex components as FortiGates, so are less likely to fail. However, for even greater reliability, a configuration is required that includes redundant connections between the cluster the networks that it is connected to.

FortiGate models that support 802.3ad Aggregate or Redundant interfaces can be used to create a cluster configuration called full mesh HA. Full mesh HA is a method of reducing the number of single points of failure on a network that includes an HA cluster.

This redundant configuration can be achieved using FortiGate 802.3ad Aggregate or Redundant interfaces and a full mesh HA configuration. In a full mesh HA configuration, you connect an HA cluster consisting of two or more FortiGates to the network using 802.3ad Aggregate or Redundant interfaces and redundant switches. Each 802.3ad Aggregate or Redundant interface is connected to two switches and both of these switches are connected to the network. In addition you must set up an IEEE 802.1Q (also called Dot1Q) or ISL link between the redundant switches connected to the Aggregate or Redundant interfaces.

The resulting full mesh configuration, an example is shown below, includes redundant connections between all network components. If any single component or any single connection fails, traffic automatically switches to the redundant component and connection and traffic flow resumes.

Single points of failure in a standalone and HA network configuration



Full mesh HA and redundant heartbeat interfaces

A full mesh HA configuration also includes redundant HA heartbeat interfaces. At least two heartbeat interfaces should be selected in the HA configuration and both sets of HA heartbeat interfaces should be connected. The HA heartbeat interfaces do not have to be configured as redundant interfaces because the FGCP handles failover between heartbeat interfaces.

Full mesh HA, redundant interfaces and 802.3ad aggregate interfaces

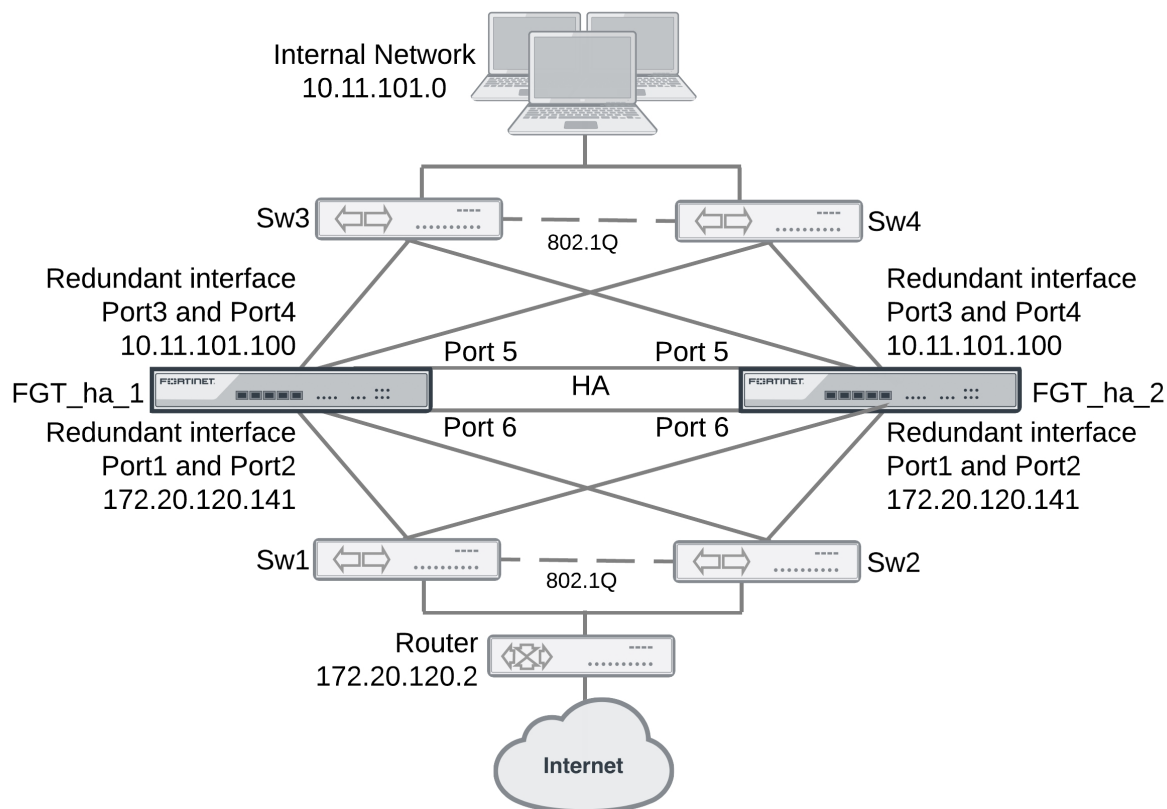
Full mesh HA is supported for both redundant interfaces and 802.3ad aggregate interfaces. In most cases you would simply use redundant interfaces. However, if your switches support 802.3ad aggregate interfaces and split multi-trunking you can use aggregate interfaces in place of redundant interfaces for full mesh HA. One advantage of using aggregate interfaces is that all of the physical interfaces in the aggregate interface can send and receive packets. As a result, using aggregate interfaces may increase the bandwidth capacity of the cluster.

Usually redundant and aggregate interfaces consist of two physical interfaces. However, you can add more than two physical interfaces to a redundant or aggregate interface. Adding more interfaces can increase redundancy protection. Adding more interfaces can also increase bandwidth capacity if you are using 802.3ad aggregate interfaces.

Example full mesh HA configuration

The following figure shows a full mesh HA configuration with a cluster of two FortiGate. This section describes the FortiGate configuration settings and network components required for a full mesh HA configuration. This section also contains example steps for setting up this full mesh HA configuration. The procedures in this section describe one of many possible sequences of steps for configuring full mesh HA. As you become more experienced with FortiOS, HA, and full mesh HA you may choose to use a different sequence of configuration steps.

Full Mesh HA configuration



For simplicity these procedures assume that you are starting with two new FortiGates set to the factory default configuration. However, starting from the default configuration is not a requirement for a successful HA deployment. FortiGate HA is flexible enough to support a successful configuration from many different starting points.

These procedures describe how to configure a cluster operating in NAT mode because NAT is the default FortiGate operating mode. However, the steps are the same if the cluster operates in transparent mode. You can either switch the cluster units to operate in transparent mode before beginning these procedures, or you can switch the cluster to operate in transparent mode after HA is configured and the cluster is connected and operating.

Full mesh HA configuration

The two FortiGates (FGT_ha_1 and FGT_ha_2) can be operating in NAT or transparent mode. Aside from the standard HA settings, the FortiGate configuration includes the following:

- The port5 and port6 interfaces configured as heartbeat interfaces. A full mesh HA configuration also includes redundant HA heartbeat interfaces.
- The port1 and port2 interfaces added to a redundant interface. Port1 is the active physical interface in this redundant interface. To make the port1 interface the active physical interface it should appear above the port2 interface in the redundant interface configuration.
- The port3 and port4 interfaces added to a redundant interface. Port3 is the active physical interface in this redundant interface. To make the port3 interface the active physical interface it should appear above the port4 interface in the redundant interface configuration.

Full mesh switch configuration

The following redundant switch configuration is required:

- Two redundant switches (Sw3 and Sw4) connected to the internal network. Establish an 802.1Q (Dot1Q) or interswitch-link (ISL) connection between them.
- Two redundant switches (Sw1 and Sw2) connected to the internet. Establish an 802.1Q (Dot1Q) or interswitch-link (ISL) connection between them.

Full mesh network connections

Make the following physical network connections for FGT_ha_1:

- Port1 to Sw1 (active)
- Port2 to Sw2 (inactive)
- Port3 to Sw3 (active)
- Port4 to Sw4 (inactive)

Make the following physical network connections for FGT_ha_2:

- Port1 to Sw2 (active)
- Port2 to Sw1 (inactive)
- Port3 to Sw4 (active)
- Port4 to Sw3 (inactive)

How packets travel from the internal network through the full mesh cluster and to the internet

If the cluster is operating in active-passive mode and FGT_ha_2 is the primary unit, all packets take the following path from the internal network to the internet:


1. From the internal network to Sw4. Sw4 is the active connection to FGT_ha_2; which is the primary unit. The primary unit receives all packets.
2. From Sw4 to the FGT_ha_2 port3 interface. Active connection between Sw4 and FGT_ha_2. Port3 is the active member of the redundant interface.
3. From FGT_ha_2 port3 to FGT_ha_2 port1. Active connection between FGT_ha_2 and Sw2. Port1 is the active member of the redundant interface.
4. From Sw2 to the external router and the internet.

Configuring full-mesh HA - GUI

Each cluster unit must have the same HA configuration.

To configure the FortiGates for HA operation

1. Register and apply licenses to the FortiGate.



2. On the **System Information** dashboard widget, beside **Host Name** select **Change**.
3. Enter a new Host Name for this FortiGate.

New Name	FGT_ha_1
-----------------	----------

4. Go to **System > HA** and change the following settings.

Mode	Active-Active	
Group Name	Rexample1.com	
Password	RHA_pass_1	
Heartbeat Interface		
	Enable	Priority
port5	Select	50
port6	Select	50

5. Select **OK**.

The FortiGate negotiates to establish an HA cluster. When you select OK you may temporarily lose connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate interfaces change to the following

virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

6. Power off the first FortiGate.
7. Repeat these steps for the second FortiGate.

Set the second FortiGate host name to:

New Name	FGT_ha_2

To connect the cluster to your network

1. Make the following physical network connections for FGT_ha_1:
 - Port1 to Sw1 (active)
 - Port2 to Sw2 (inactive)
 - Port3 to Sw3 (active)
 - Port4 to Sw4 (inactive)
2. Make the following physical network connections for FGT_ha_2:
 - Port1 to Sw2 (active)
 - Port2 to Sw1 (inactive)
 - Port3 to Sw4 (active)
 - Port4 to Sw3 (inactive)
3. Connect Sw3 and Sw4 to the internal network.
4. Connect Sw1 and Sw2 to the external router.
5. Enable 802.1Q (Dot1Q) or ISL communication between Sw1 and Sw2 and between Sw3 and Sw4.
6. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention.

When negotiation is complete the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view the cluster dashboard and cluster members list to confirm that the cluster units are operating as a cluster.

1. View the system dashboard.

The System Information dashboard widget shows the **Cluster Name** (Rexample1.com) and the host names and serial numbers of the **Cluster Members**. The Unit Operation widget shows multiple cluster units.
2. Go to **System > HA** to view the cluster members list.

The list shows two cluster units, their host names, their roles in the cluster, and their priorities. You can use this list to confirm that the cluster is operating normally.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard does not display information for both cluster units the FortiGates are not functioning as a cluster. See [Example full mesh HA configuration on page 1475](#) to troubleshoot the cluster.

To add basic configuration settings and the redundant interfaces

Use the following steps to add a few basic configuration settings.

1. Log into the cluster GUI.
2. Go to **System > Administrators**.
3. Edit **admin** and select **Change Password**.
4. Enter and confirm a new password.

5. Select **OK**.
6. Go to **Network > Static Routes** and temporarily delete the default route.
You cannot add an interface to a redundant interface if any settings (such as the default route) are configured for it.
7. Go to **Network > Interfaces** and select **Create New > Interface** and configure the redundant interface to connect to the internet.

Name	Port1_Port2
Type	Redundant
Physical Interface Members	
Selected Interfaces	port1, port2
IP/Netmask	172.20.120.141/24

8. Select **OK**.
9. Select **Create New** and configure the redundant interface to connect to the internal network.

Name	Port3_Port4
Type	Redundant
Physical Interface Members	
Selected Interfaces	port3, port4
IP/Netmask	10.11.101.100/24
Administrative Access	HTTPS, PING, SSH

10. Select **OK**.
The virtual MAC addresses of the FortiGate interfaces change to the following. Notice that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)

- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

11. Go to **Router > Static > Static Routes**.

12. Add the default route.

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.2
Device	Port1_Port2
Distance	10

13. Select **OK**.

To configure HA port monitoring for the redundant interfaces

1. Go to **System > HA**.
2. In the cluster members list, edit the primary unit.
3. Enable **interface monitoring** the **Port1_Port2** and the **Port3_Port4** interfaces
4. Select **OK**.

Configuring full mesh HA - CLI

Each cluster must have the same HA configuration. Use the following procedure to configure the FortiGates for HA operation.

To configure the FortiGates for HA operation

1. Register and apply licenses to the FortiGate.
2. Enter a new Host Name for this FortiGate.

```
config system global
    set hostname FGT_ha_1
end
```

3. Configure HA settings.

```
config system ha
    set mode a-a
    set group-name Rexample1.com
    set password RHA_pass_1
    set hbdev port5 50 port6 50
end
```

The FortiGate negotiates to establish an HA cluster. When you select OK you may temporarily lose

connectivity with the FortiGate as the HA cluster negotiates and the FGCP changes the MAC address of the FortiGate interfaces. The MAC addresses of the FortiGate interfaces change to the following virtual MAC addresses:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01
- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-0b
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0e
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

To be able to reconnect sooner, you can update the ARP table of your management PC by deleting the ARP table entry for the FortiGate (or just deleting all arp table entries). You may be able to delete the arp table of your management PC from a command prompt using a command similar to `arp -d`.

You can use the `get hardware nic` (or `diagnose hardware deviceinfo nic`) CLI command to view the virtual MAC address of any FortiGate interface. For example, use the following command to view the port1 interface virtual MAC address (`Current_HWaddr`) and the port1 permanent MAC address (`Permanent_HWaddr`):

```
get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:00
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

4. Power off the first FortiGate.
5. Repeat these steps for the second FortiGate.

Set the other FortiGate host name to:

```
config system global
    set hostname FGT_ha_2
end
```

To connect the cluster to your network

1. Make the following physical network connections for FGT_ha_1:
 - Port1 to Sw1 (active)
 - Port2 to Sw2 (inactive)
 - Port3 to Sw3 (active)
 - Port4 to Sw4 (inactive)
2. Make the following physical network connections for FGT_ha_2:
 - Port1 to Sw2 (active)
 - Port2 to Sw1 (inactive)
 - Port3 to Sw4 (active)
 - Port4 to Sw3 (inactive)
3. Connect Sw3 and Sw4 to the internal network.
4. Connect Sw1 and Sw2 to the external router.
5. Enable 802.1Q (Dot1Q) or ISL communication between Sw1 and Sw2 and between Sw3 and Sw4.
6. Power on the cluster units.

The units start and negotiate to choose the primary unit and the subordinate unit. This negotiation occurs with no user intervention.

When negotiation is complete the cluster is ready to be configured for your network.

To view cluster status

Use the following steps to view cluster status from the CLI.

1. Log into the CLI.
2. Enter `get system status` to verify the HA status of the cluster unit that you logged into.

If the command output includes `Current HA mode: a-a, master`, the cluster units are operating as a cluster and you have connected to the primary unit.

If the command output includes `Current HA mode: a-a, backup`, you have connected to a subordinate unit.

If the command output includes `Current HA mode: standalone` the cluster unit is not operating in HA mode.

3. Enter the following command to confirm the HA configuration of the cluster:

```
get system ha status
HA Health Status: OK
Model: FortiGate-XXXX
Mode: HA A-P
Group: 0
Debug: 0
Cluster Uptime: 7 days 00:30:26
.
.
.
```

You can use this command to confirm that the cluster is healthy and operating normally, some information about the cluster configuration, and information about how long the cluster has been

operating. Information not shown in this example includes how the primary unit was selected, configuration synchronization status, usage stats for each cluster unit, heartbeat status, and the relative priorities of the cluster units.

4. Use the `execute ha manage` command to connect to the other cluster unit's CLI and use these commands to verify cluster status.

To troubleshoot the cluster configuration

If the cluster members list and the dashboard does not display information for both cluster units the FortiGates are not functioning as a cluster. See [Example full mesh HA configuration on page 1475](#) to troubleshoot the cluster.

To add basic configuration settings and the redundant interfaces

Use the following steps to add a few basic configuration settings. Some steps use the CLI and some the GUI.

1. Log into the cluster CLI.
2. Add a password for the admin administrative account.

```
config system admin
  edit admin
    set password <password_str>
  end
```

3. Temporarily delete the default route.

You cannot add an interface to a redundant interface if any settings (such as the default route) are configured for it.

```
config router static
  delete 1
end
```

4. Go to **System > Network > Interface** and select **Create New** to add the redundant interface to connect to the internet.
5. Add the redundant interface to connect to the internet.

```
config system interface
  edit Port1_Port2
    set type redundant
    set member port1 port2
  end
```

6. Add the redundant interface to connect to the internal network.

```
config system interface
  edit Port3_Port4
    set type redundant
    set member port3 port4
  end
```

The virtual MAC addresses of the FortiGate interfaces change to the following. Note that port1 and port2 both have the port1 virtual MAC address and port3 and port4 both have the port3 virtual MAC address:

- port1 interface virtual MAC: 00-09-0f-09-00-00
- port10 interface virtual MAC: 00-09-0f-09-00-01

- port11 interface virtual MAC: 00-09-0f-09-00-02
- port12 interface virtual MAC: 00-09-0f-09-00-03
- port13 interface virtual MAC: 00-09-0f-09-00-04
- port14 interface virtual MAC: 00-09-0f-09-00-05
- port15 interface virtual MAC: 00-09-0f-09-00-06
- port16 interface virtual MAC: 00-09-0f-09-00-07
- port17 interface virtual MAC: 00-09-0f-09-00-08
- port18 interface virtual MAC: 00-09-0f-09-00-09
- port19 interface virtual MAC: 00-09-0f-09-00-0a
- port2 interface virtual MAC: 00-09-0f-09-00-00 (same as port1)
- port20 interface virtual MAC: 00-09-0f-09-00-0c
- port3 interface virtual MAC: 00-09-0f-09-00-0d
- port4 interface virtual MAC: 00-09-0f-09-00-0d (same as port3)
- port5 interface virtual MAC: 00-09-0f-09-00-0f
- port6 interface virtual MAC: 00-09-0f-09-00-10
- port7 interface virtual MAC: 00-09-0f-09-00-11
- port8 interface virtual MAC: 00-09-0f-09-00-12
- port9 interface virtual MAC: 00-09-0f-09-00-13

7. Go to Router > Static > Static Routes.

8. Add the default route.

```
config router static
edit 1
set dst 0.0.0.0 0.0.0.0
set gateway 172.20.120.2
set device Port1_Port2
end
```

To configure HA port monitoring for the redundant interfaces

1. Enter the following command to configure port monitoring for the redundant interfaces:

```
config system ha
set monitor Port1_Port2 Port3_Port4
end
```

Troubleshooting full mesh HA

Troubleshooting full mesh HA clusters is similar to troubleshooting any cluster (see [FGCP configuration examples and troubleshooting on page 1412](#) or [Virtual clusters on page 1463](#)). The configuration and operation of a full mesh HA cluster is very similar to the configuration and operation of a standard cluster. The only differences relate to the configuration, connection, and operation of the redundant interfaces and redundant switches.

- Make sure the redundant interfaces and switches are connected correctly. With so many connections it is possible to make mistakes or for cables to become disconnected.
- Confirm that the configuration of the cluster unit 802.3ad Aggregate or Redundant interfaces is correct according to the configuration procedures in this chapter.
- In some configurations with some switch hardware, MAC-learning delays on the inter-switch links on the surrounding topologies may occur. The delays occur if the gratuitous ARP packets sent by the cluster after a failover

are delayed by the switches before being sent across the inter-switch link. If this happens the surrounding topologies may be delayed in recognizing the failover and will keep sending packets to the MAC address of the failed primary unit resulting in lost traffic. Resolving this problem may require changing the configuration of the switch or replacing them with switch hardware that does not delay the gratuitous ARP packets.

Operating clusters and virtual clusters

With some exceptions, you can operate a cluster in much the same way as you operate a standalone FortiGate. This chapter describes those exceptions and also the similarities involved in operating a cluster instead of a standalone FortiGate.

Operating a cluster

The configurations of all of the FortiGates in a cluster are synchronized so that the cluster units can simulate a single FortiGate. Because of this synchronization, you manage the HA cluster instead of managing the individual cluster units. You manage the cluster by connecting to the GUI using any cluster interface configured for HTTPS or HTTP administrative access. You can also manage the cluster by connecting to the CLI using any cluster interface configured for SSH or telnet administrative access.

The cluster GUI dashboard displays the cluster name, the host name and serial number of each cluster member, and also shows the role of each unit in the cluster. The roles can be master (primary unit) and slave (subordinate units). The dashboard also displays a cluster unit front panel illustration.

You can also go to **System > HA** to view the cluster members list. This includes status information for each cluster unit. You can also use the cluster members list for a number of cluster management functions including changing the HA configuration of an operating cluster, changing the host name and device priority of a subordinate unit, and disconnecting a cluster unit from a cluster. See [Cluster members list on page 1505](#).

You can use log messages to view information about the status of the cluster. See [Clusters and logging on page 1497](#).

You can use SNMP to manage the cluster by configuring a cluster interface for SNMP administrative access. Using an SNMP manager you can get cluster configuration information and receive traps. See [Clusters and SNMP on page 1500](#).

You can configure a reserved management interface to manage individual cluster units. You can use this interface to access the GUI or CLI and to configure SNMP management for individual cluster units. See [Managing individual cluster units using a reserved out-of-band management interface on page 1488](#).

You can manage individual cluster units by using SSH, telnet, or the CLI console on the GUI dashboard to connect to the CLI of the cluster. From the CLI you can use the `execute ha manage` command to connect to the CLI of any unit in the cluster.

You can also manage individual cluster units by using a null-modem cable to connect to any cluster unit CLI. From there you can use the `execute ha manage` command to connect to the CLI of each unit in the cluster.

Operating a virtual cluster

Managing a virtual cluster is very similar to managing a cluster that does not contain multiple virtual domains. Most of the information in this chapter applies to managing both kinds of clusters. This section describes what is different when managing a virtual cluster.

If virtual domains are enabled, the cluster GUI dashboard displays the cluster name and the role of each cluster unit in virtual cluster 1 and virtual cluster 2.

The configuration and maintenance options that you have when you connect to a virtual cluster GUI or CLI depend on the virtual domain that you connect to and the administrator account that you use to connect.

If you connect to a cluster as the administrator of a virtual domain, you connect directly to the virtual domain. Since HA virtual clustering is a global configuration, virtual domain administrators cannot see HA configuration options. However, virtual domain administrators see the host name of the cluster unit that they are connecting to on the web browser title bar or CLI prompt. This host name is the host name of the primary unit for the virtual domain. Also, when viewing log messages the virtual domain administrator can select to view log messages for either of the cluster units.

If you connect to a virtual cluster as the admin administrator you connect to the global GUI or CLI. Even so, you are connecting to an interface and to the virtual domain that the interface has been added to. The virtual domain that you connect to does not make a difference for most configuration and maintenance operations. However, there are a few exceptions. You connect to the FortiGate that functions as the primary unit for the virtual domain. So the host name displayed on the web browser title bar and on the CLI is the host name of this primary unit.

Managing individual cluster units using a reserved out-of-band management interface

You can provide direct management access to all cluster units by reserving up to four management interfaces as part of the HA configuration. Once a management interface is reserved, you can configure a different IP address, administrative access and other interface settings for each management interface for each cluster unit. Then by connecting these interfaces of each cluster unit to your network you can manage each cluster unit separately from different IP addresses. Configuration changes to the reserved management interfaces are not synchronized to other cluster units.



You can also configure an in-band management interface for a cluster unit. See ["Managing individual cluster units using an in-band management IP address" on page 1494](#).

Reserved management interfaces provide direct management access to each cluster unit and give each cluster unit a different identity on your network. This simplifies using external services, such as SNMP, to separately monitor and manage each cluster unit.



The reserved management interfaces are not assigned HA virtual MAC addresses like other cluster interfaces. Instead the reserved management interfaces retain the permanent hardware address of the physical interface unless you change it using the `config system interface` command.

Reserved management interfaces and their IP addresses should not be used for managing a cluster using FortiManager. To correctly manage a FortiGate HA cluster with FortiManager use the IP address of one of the cluster unit interfaces.

If you enable SNMP administrative access for a reserved management interface you can use SNMP to monitor each cluster unit using the reserved management interface IP address. To monitor each cluster unit using SNMP, just add the IP address of each cluster unit's reserved management interface to the SNMP server configuration. You must also enable direct management of cluster members in the cluster SNMP configuration.

If you enable HTTPS or HTTP administrative access for the reserved management interfaces you can connect to the GUI of each cluster unit. Any configuration changes made to any of the cluster units is automatically synchronized to all cluster units. From the subordinate units the GUI has the same features as the primary unit except that unit-specific information is displayed for the subordinate unit, for example:

- The **Dashboard System Information** widget displays the subordinate unit serial number but also displays the same information about the cluster as the primary unit

- On the Cluster members list (go to **System > HA**) you can change the HA configuration of the subordinate unit that you are logged into. For the primary unit and other subordinate units you can change only the host name and device priority.
- Log Access displays the logs of the subordinate that you are logged into first. You use the HA Cluster list to view the log messages of other cluster units including the primary unit.

If you enable SSH or TELNET administrative access for the reserved management interfaces you can connect to the CLI of each cluster unit. The CLI prompt contains the host name of the cluster unit that you have connected to. Any configuration changes made to any of the cluster units is automatically synchronized to all cluster units. You can also use the `execute ha manage` command to connect to other cluster unit CLIs.

The reserved management interface is available in NAT and in transparent mode. It is also available if the cluster is operating with multiple VDOMs. In transparent mode you cannot normally add an IP address to an interface. However, you can add an IP address to the reserved management interface.

Using an HA reserved management interface for FortiSandbox, SNMP and other management services

By default, management services such as SNMP, remote logging, remote authentication and communication with FortiSandbox and so on use a cluster interface. As a result communication from each cluster unit comes from a cluster interface instead of from the interface of an individual cluster unit and not from the HA reserved management interface.

If you want to use a HA reserved management interface for these features you must enter the following command:

```
config system ha
    set ha-direct enable
end
```

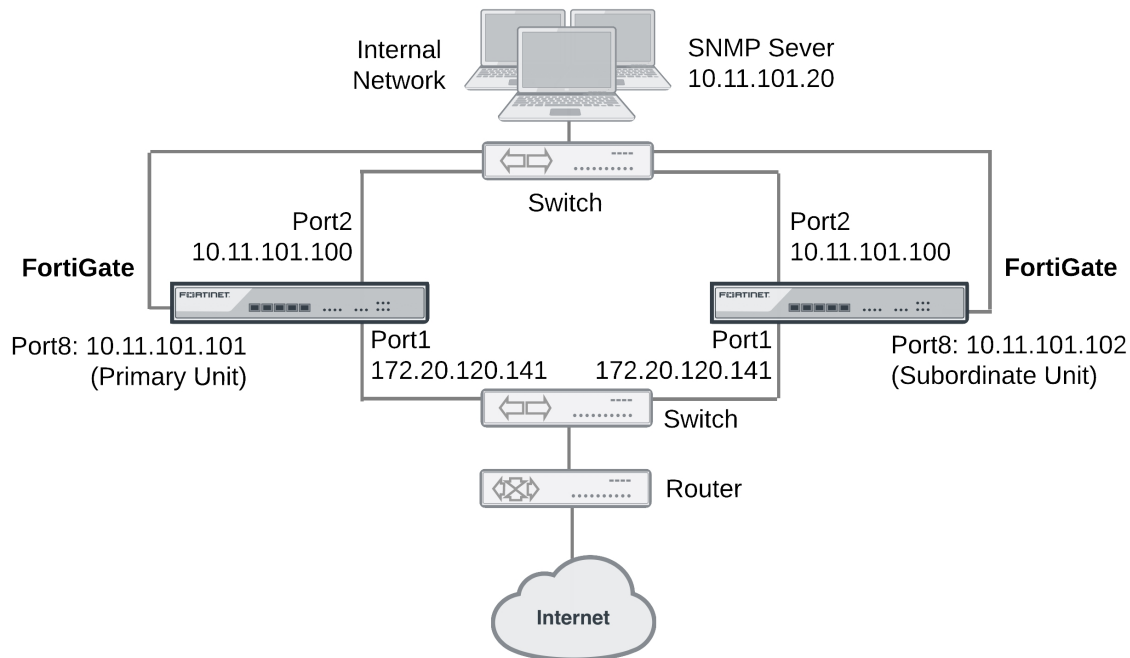
The result is that all management services can use the HA reserved management interfaces. This means that individual cluster units send log messages and communicate with FortiSandbox and so on using the HA reserved management interfaces instead of one of the cluster interfaces. This allows you to manage each cluster unit separately and to separate the management traffic from each cluster unit. This can also be useful if each cluster unit is in a different location.

If you just want to use HA reserved management interfaces for SNMP remote management you can enable `ha-direct` in the SNMP configuration as shown in the following example.

Configuring the reserved management interface and SNMP remote management of individual cluster units

This example describes how to configure SNMP remote management of individual cluster units using an HA reserved management interface. The configuration consists of two FortiGate-620B units already operating as a cluster. In the example, the port8 interface of each cluster unit is connected to the internal network using the switch and configured as the reserved management interface.

SNMP remote management of individual cluster units



To configure the reserved management interface - GUI

From the GUI you can also configure IPv4 and IPv6 default routes that are only used by the reserved management interface.

1. Go to **System > HA**.
2. Edit the primary unit.
3. Select **Management Interface Reservation** and select port8.
4. Set **Gateway** to 10.11.101.2.
5. Set **IPv6 Gateway** to 2001:db8:0:2::20
6. Select **OK**.

To configure the reserved management interface - CLI

From the CLI you can also configure IPv4 and IPv6 default routes that are only used by the reserved management interface.

1. Log into the CLI of any cluster unit.
2. Enter the following command to enable the reserved management interface, set port8 as the reserved interface, and add an IPv4 default route of 10.11.101.2 and an IPv6 default route of 2001:db8:0:2::20 for the reserved management interface.

```
config system ha
  set ha-mgmt-status enable
  config ha-mgmt-interfaces
    edit 1
      set interface port8
```

```

set gateway 10.11.101.2
set gateway6 2001:db8:0:2::20
end

```

The reserved management interface default route is not synchronized to other cluster units.

To change the primary unit reserved management interface configuration - GUI

You can change the IP address of the primary unit reserved management interface from the primary unit GUI. Configuration changes to the reserved management interface are not synchronized to other cluster units.

1. From a PC on the internal network, browse to <http://10.11.101.100> and log into the cluster GUI.

This logs you into the primary unit GUI.

You can identify the primary unit from its serial number or host name that appears on the System Information dashboard widget.

2. Go to **Network > Interfaces** and edit the port8 interface as follows:

Alias	primary_reserved
IP/Netmask	10.11.101.101/24
Administrative Access	Ping, SSH, HTTPS, SNMP

3. Select **OK**.

You can now log into the primary unit GUI by browsing to <https://10.11.101.101>. You can also log into this primary unit CLI by using an SSH client to connect to 10.11.101.101.

To change subordinate unit reserved management interface configuration - CLI

At this point you cannot connect to the subordinate unit reserved management interface because it does not have an IP address. Instead, this procedure describes connecting to the primary unit CLI and using the `execute ha manage` command to connect to subordinate unit CLI to change the port8 interface. You can also use a serial connection to the cluster unit CLI. Configuration changes to the reserved management interface are not synchronized to other cluster units.

1. Connect to the primary unit CLI and use the `execute ha manage` command to connect to a subordinate unit CLI.

You can identify the subordinate unit from its serial number or host name. The host name appears in the CLI prompt.

2. Enter the following command to change the port8 IP address to 10.11.101.102 and set management access to HTTPS, ping, SSH, and SNMP.

```

config system interface
  edit port8
    set ip 10.11.101.102/24
    set allowaccess https ping ssh snmp
  end

```

You can now log into the subordinate unit GUI by browsing to <https://10.11.101.102>. You can also log into this subordinate unit CLI by using an SSH client to connect to 10.11.101.102.

To configure the cluster for SNMP management using the reserved management interfaces - CLI

This procedure describes how to configure the cluster to allow the SNMP server to get status information from the primary unit and the subordinate unit. The SNMP configuration is synchronized to all cluster units. To support using the reserved management interfaces, you must add at least one HA direct management host to an SNMP community. If your SNMP configuration includes SNMP users with user names and passwords you must also enable HA direct management for SNMP users.

1. Enter the following command to add an SNMP community called `Community` and add a host to the community for the reserved management interface of each cluster unit. The host includes the IP address of the SNMP server (10.11.101.20).

```
config system snmp community
edit 1
set name Community
config hosts
edit 1
set ha-direct enable
set ip 10.11.101.20
end
end
```



Enabling ha-direct in non-HA environments makes SNMP unusable.

- 2.
3. Enter the following command to add an SNMP user for the reserved management interface.

```
config system snmp user
edit 1
set ha-direct enable
set notify-hosts 10.11.101.20
end
```

Configure other settings as required.

To get CPU, memory, and network usage of each cluster unit using the reserved management IP addresses

From the command line of an SNMP manager, you can use the following SNMP commands to get CPU, memory and network usage information for each cluster unit. In the examples, the community name is `Community`. The commands use the MIB field names and OIDs listed below.

Enter the following commands to get CPU, memory and network usage information for the primary unit with reserved management IP address 10.11.101.101 using the MIB fields:

```
snmpget -v2c -c Community 10.11.101.101 fgHaStatsCpuUsage
snmpget -v2c -c Community 10.11.101.101 fgHaStatsMemUsage
snmpget -v2c -c Community 10.11.101.101 fgHaStatsNetUsage
```

Enter the following commands to get CPU, memory and network usage information for the primary unit with reserved management IP address 10.11.101.101 using the OIDs:

```
snmpget -v2c -c Community 10.11.101.101 1.3.6.1.4.1.12356.101.13.2.1.1.3.1
snmpget -v2c -c Community 10.11.101.101 1.3.6.1.4.1.12356.101.13.2.1.1.4.1
snmpget -v2c -c Community 10.11.101.101 1.3.6.1.4.1.12356.101.13.2.1.1.5.1
```

Enter the following commands to get CPU, memory and network usage information for the subordinate unit with reserved management IP address 10.11.101.102 using the MIB fields:

```
snmpget -v2c -c Community 10.11.101.102 fgHaStatsCpuUsage
snmpget -v2c -c Community 10.11.101.102 fgHaStatsMemUsage
snmpget -v2c -c Community 10.11.101.102 fgHaStatsNetUsage
```

Enter the following commands to get CPU, memory and network usage information for the subordinate unit with reserved management IP address 10.11.101.102 using the OIDs:

```
snmpget -v2c -c Community 10.11.101.102 1.3.6.1.4.1.12356.101.13.2.1.1.3.1
snmpget -v2c -c Community 10.11.101.102 1.3.6.1.4.1.12356.101.13.2.1.1.4.1
snmpget -v2c -c Community 10.11.101.102 1.3.6.1.4.1.12356.101.13.2.1.1.5.1
```

Adding firewall local-in policies for the dedicated HA management interface

To add local-in policies for the dedicated management interface, enable `ha-mgmt-intf-only` and set `intf` to `any`. Enabling `ha-mgmt-intf-only` means the local-in policy applies only to the VDOM that contains the dedicated HA management interface. For example:

```
config firewall local-in-policy
edit 0
set ha-mgmt-intf-only enable
set intf any
set srcaddr internal-net
set dstaddr mgmt-int
set action accept
set service HTTPS
set schedule weekdays
end
```

NTP over dedicated HA management interfaces

If you set up dedicated management interfaces on each cluster unit, if NTP is enabled, the primary unit contacts an NTP server using the dedicated management interface. System time is then synchronized to the backup units through the HA heartbeat.

Example CLI:

```
config system interface
edit port5
set ip 172.16.79.46 255.255.255.0
end

config system ha
set group-name FGT-HA
set mode a-p
set ha-mgmt-status enable
config ha-mgmt-interfaces
edit 1
set interface port5
set gateway 172.16.79.1
end
set ha-direct enable
end

config system ntp
set ntpsync enable
set syncinterval 5
```

```
end
```

Managing individual cluster units using an in-band management IP address

You can use the following command to add an in-band management IP address to an individual cluster unit interface that is also connected to a network and processing traffic. The in-band management IP address is an alternative to the reserved HA management interface feature and does not require reserving an interface just for management access.

```
config system interface
  edit port1
    set management-ip 172.20.121.155/24
  end
```

The management IP address is accessible from the network that the cluster interface is connected to. This setting is not synchronized so each cluster unit can have their own in-band management IP addresses. You can add a management IP address to one or more interfaces of each cluster unit.

The in-band management IP address should be on the same subnet as the interface you are adding it to, but cannot be on the same subnet as other interface IP addresses.

You can connect to the in-band management IP address using the interface's administrative access settings. The in-band management IP only supports the following subset of administrative access settings: ping, Telnet, HTTP, HTTPS, and SNMP.

For example, use the following command to add an in-band management IP address and allow access using HTTPS, SSH and SNMP:

```
config system interface
  edit port23
    set management-ip 172.25.12.5/24
    set allowaccess https ssh snmp
  end
```

Managing individual cluster units in a virtual cluster

You can select the HA option **Do NOT Synchronize Management VDOM Configuration** if you have enabled multiple VDOMS and set a VDOM other than the root VDOM to be the management VDOM. You can select this option to prevent the management VDOM configuration from being synchronized between cluster units in a virtual cluster. This allows you to add an interface to the VDOM in each cluster unit and then to give the interfaces different IP addresses in each cluster unit, allowing you to manage each cluster unit separately.

You can also enable this feature using the following command:

```
config system ha
  set standalone-mgmt-vdom enable
end
```



This feature must be disabled to manage a cluster using FortiManager.

Shutting down or rebooting the primary unit

You can shutdown or reboot the primary unit from the primary unit GUI by selecting **Shutdown** or **Reboot** from the **Admin** menu. From the primary unit CLI you can use the `execute reboot` or `execute shutdown` commands to shutdown or reboot the primary unit.

During the shutdown the primary unit first becomes the backup unit before shutting down allowing the backup unit to become the new primary unit and avoiding a split brain scenario. This behavior only happens when you manually shutdown or reboot the primary unit.

The primary unit acts as a router for subordinate unit management traffic

HA uses routing and inter-VDOM links to route subordinate unit management traffic through the primary unit to the network. Similar to a standalone FortiGate, subordinate units may generate their own management traffic, including:

- DNS queries.
- FortiGuard Web Filtering rating requests.
- Log messages to be sent to a FortiAnalyzer unit, to a syslog server, or to the FortiGuard Analysis and Management Service.
- Log file uploads to a FortiAnalyzer unit.
- Quarantine file uploads to a FortiAnalyzer unit.
- SNMP traps.
- Communication with remote authentication servers (RADIUS, LDAP, TACACS+ and so on)


Subordinate units send this management traffic over the HA heartbeat link to the primary unit. The primary unit forwards the management traffic to its destination. The primary unit also routes replies back to the subordinate unit in the same way.

HA uses a hidden VDOM called `vsys_ha` for HA operations. The `vsys_ha` VDOM includes the HA heartbeat interfaces, and all communication over the HA heartbeat link goes through the `vsys_ha` VDOM. To provide communication from a subordinate unit to the network, HA adds hidden inter-VDOM links between the primary unit management VDOM and the primary unit `vsys_ha` VDOM. By default, `root` is the management VDOM.

Management traffic from the subordinate unit originates in the subordinate unit `vsys_ha` VDOM. The `vsys_ha` VDOM routes the management traffic over the HA heartbeat link to the primary unit `vsys_ha` VDOM. This management traffic is then routed to the primary unit management VDOM and from there out onto the network.

DNS queries and FortiGuard Web Filtering and Email Filter requests are still handled by the HA proxy so the primary unit and subordinate units share the same DNS query cache and the same FortiGuard Web Filtering and Email Filter cache. In a virtual clustering configuration, the cluster unit that is the primary unit for the management virtual domain maintains the FortiGuard Web Filtering, Email Filtering, and DNS query cache.

Subordinate unit management traffic path



Cluster communication with RADIUS and LDAP servers

In an active-passive cluster, only the primary unit processes traffic, so the primary unit communicates with RADIUS or LDAP servers. In a cluster that is operating in active-active mode, subordinate units send RADIUS and LDAP requests to the primary unit over the HA heartbeat link and the primary unit routes them to their destination. The primary unit relays the responses back to the subordinate unit.

Clusters and FortiGuard services

This section describes how various FortiGate HA clustering configurations communicate with the FDN.

In an operating cluster, the primary unit communicates directly with the FortiGuard Distribution Network (FDN). Subordinate units also communicate directly with the FDN but as described above, all communication between subordinate units and the FDN is routed through the primary unit.

You must register and license all of the units in a cluster for all required FortiGuard services, both because all cluster units communicate with the FDN and because any cluster unit could potentially become the primary unit.

FortiGuard and active-passive clusters

For an active-passive cluster, only the primary unit processes traffic. Even so, all cluster units communicate with the FDN. Only the primary unit sends FortiGuard Web Filtering and Antispam requests to the FDN. All cluster units receive FortiGuard Antivirus, IPS, and application control updates from the FDN.

In an active-passive cluster the FortiGuard Web Filter and Email Filter caches are located on the primary unit in the same way as for a standalone FortiGate. The caches are not shared among cluster units so after a failover the new primary unit must build up new caches.

In an active-passive cluster all cluster units also communicate with the FortiGuard Analysis and Management Service (FAMS).

FortiGuard and active-active clusters

For an active-active cluster, both the primary unit and the subordinate units process traffic. Communication between the cluster units and the FDN is the same as for active-passive clusters with the following exception.

Because the subordinate units process traffic, they may also be making FortiGuard Web Filtering and Email Filter requests. The primary unit receives all such requests from the subordinate units and relays them to the FDN and then relays the FDN responses back to the subordinate units. The FortiGuard Web Filtering and Email Filtering

URL caches are maintained on the primary unit. The primary unit caches are used for primary and subordinate unit requests.

FortiGuard and virtual clustering

For a virtual clustering configuration the management virtual domain of each cluster unit communicates with the FDN. The cluster unit that is the primary unit for the management virtual domain maintains the FortiGuard Web Filtering and Email Filtering caches. All FortiGuard Web Filtering and Email Filtering requests are proxied by the management VDOM of the cluster unit that is the primary unit for the management virtual domain.

Clusters and logging

This section describes the log messages that provide information about how HA is functioning, how to view and manage logs for each unit in a cluster, and provides some example log messages that are recorded during specific cluster events.

You configure logging for a cluster in the same way as you configuring logging for a standalone FortiGate. Log configuration changes made to the cluster are synchronized to all cluster units.

All cluster units record log messages separately to the individual cluster unit's log disk, to the cluster unit's system memory, or both. You can view and manage log messages for each cluster unit from the cluster GUI Log Access page.

When remote logging is configured, all cluster units send log messages to remote FortiAnalyzer units or other remote servers as configured. HA uses routing and inter-VDOM links to route subordinate unit log traffic through the primary unit to the network.

When you configure a FortiAnalyzer unit to receive log messages from a FortiGate cluster, you should add a cluster to the FortiAnalyzer unit configuration so that the FortiAnalyzer unit can receive log messages from all cluster units.

Viewing and managing log messages for individual cluster units

This section describes how to view and manage log messages for an individual cluster unit.

To view HA cluster log messages

1. Log into the cluster GUI.
2. Go to **Log&Report > Log Config > Log Settings > GUI Preferences** and select to display logs from **Memory, Disk or FortiAnalyzer**.

For each log display, the **HA Cluster** list displays the serial number of the cluster unit for which log messages are displayed. The serial numbers are displayed in order in the list.

3. Set **HA Cluster** to the serial number of one of the cluster units to display log messages for that unit.

About HA event log messages

HA event log messages always include the host name and serial number of the cluster unit that recorded the message. HA event log messages also include the HA state of the unit and also indicate when a cluster unit switches (or moves) from one HA state to another. Cluster units can operate in the HA states listed below:

HA states

Hello	A FortiGate configured for HA operation has started up and is looking for other FortiGates with which to form a cluster.
Work	In an active-passive cluster a cluster unit is operating as the primary unit. In an active-active cluster unit is operating as the primary unit or a subordinate unit.
Standby	In an active-passive cluster the cluster unit is operating as a subordinate unit.

HA log Event log messages also indicate the virtual cluster that the cluster unit is operating in as well as the member number of the unit in the cluster. If virtual domains are not enabled, all clusters unit are always operating in virtual cluster 1. If virtual domains are enabled, a cluster unit may be operating in virtual cluster 1 or virtual cluster 2. The member number indicates the position of the cluster unit in the cluster members list. Member 0 is the primary unit. Member 1 is the first subordinate unit, member 2 is the second subordinate unit, and so on.

HA log messages

See the FortiOS log message reference for a listing of and descriptions of the HA log messages.

FortiGate HA message "HA master heartbeat interface <intf_name> lost neighbor information"

The following HA log messages may be recorded by an operating cluster:

```
2009-02-16 11:06:34 device_id=FG2001111111 log_id=0105035001 type=event subtype=ha
pri=critical vd=root msg="HA slave heartbeat interface internal lost neighbor information"
```

```
2009-02-16 11:06:40 device_id=FG2001111111 log_id=0105035001 type=event subtype=ha
pri=notice vd=root msg="Virtual cluster 1 of group 0 detected new joined HA member"
```

```
2009-02-16 11:06:40 device_id=FG2001111111 log_id=0105035001 type=event subtype=ha
pri=notice vd=root msg="HA master heartbeat interface internal get peer information"
```

These log messages indicate that the cluster units could not connect to each other over the HA heartbeat link for the period of time that is given by hb-interval x hb-lost-threshold, which is 1.2 seconds with the default values.

To diagnose this problem

1. Check all heartbeat interface connections including cables and switches to make sure they are connected and operating normally.
2. Use the following commands to display the status of the heartbeat interfaces.


```
get hardware nic <heartbeat_interface_name>
diagnose hardware deviceinfo nic <heartbeat_interface_name>
```

The status information may indicate the interface status and link status and also indicate if a large number of errors have been detected.
3. If the log message only appears during peak traffic times, increase the tolerance for missed HA heartbeat packets by using the following commands to increase the lost heartbeat threshold and heartbeat interval:

```
config system ha
  set hb-lost-threshold 12
  set hb-interval 4
end
```

These settings multiply by 4 the loss detection interval. You can use higher values as well.

This condition can also occur if the cluster units are located in different buildings or even different geographical locations. Called a distributed cluster, as a result of the separation it may take a relatively long time for heartbeat packets to be transmitted between cluster units. You can support a distributed cluster by increasing the heartbeat interval so that the cluster expects extra time between heartbeat packets.

4. Optionally disable session-pickup to reduce the processing load on the heartbeat interfaces.
5. Instead of disabling session-pickup you can enable `session-pickup-delay` to reduce the number of sessions that are synchronized. With this option enabled only sessions that are active for more than 30 seconds are synchronized.

It may be useful to monitor CPU and memory usage to check for low memory and high CPU usage. You can configure event logging to monitor CPU and memory usage. You can also enable the CPU over usage and memory low SNMP events.

Once this monitoring is in place, try and determine if there have been any changes in the network or an increase of traffic recently that could be the cause. Check to see if the problem happens frequently and if so what the pattern is.

To monitor the CPU of the cluster units and troubleshoot further, use the following procedure and commands:

```
get system performance status
get system performance top 2
diagnose sys top 2
```

These commands repeated at frequent intervals will show the activity of the CPU and the number of sessions.

Search the Fortinet Knowledge Base for articles about monitoring CPU and Memory usage.

If the problem persists, gather the following information (a console connection might be necessary if connectivity is lost) and provide it to Technical Support when opening a ticket:

- Debug log from the GUI: **System > Advanced > Download Debug Log**
- CLI command output:

```
diagnose sys top 2 (keep it running for 20 seconds)
get system performance status (repeat this command multiple times to get good samples)
get system ha status
diagnose sys ha status
diagnose sys ha dump-by {all options}
diagnose netlink device list
diagnose hardware deviceinfo nic <heartbeat-interface-name>
execute log filter category 1
execute log display
```

Formatting cluster unit hard disks (log disks)

If you need to format the hard disk (also called log disk or disk storage) of one or more cluster units you should disconnect the unit from the cluster and use the `execute formatlogdisk` command to format the cluster unit hard disk then add the unit back to the cluster.

For information about how to remove a unit from a cluster and add it back, see [Disconnecting a cluster unit from a cluster on page 1518](#) and [Adding a disconnected FortiGate back to its cluster on page 1519](#).

Once you add the cluster unit with the formatted log disk back to the cluster you should make it the primary unit before removing other units from the cluster to format their log disks and then add them back to the cluster.

Clusters and SNMP

You can use SNMP to manage a cluster by configuring a cluster interface for SNMP administrative access. Using an SNMP manager you can get cluster configuration and status information and receive traps.

You configure SNMP for a cluster in the same way as configuring SNMP for a standalone FortiGate. SNMP configuration changes made to the cluster are shared by all cluster units.

Each cluster unit sends its own traps and SNMP manager systems can use SNMP get commands to query each cluster unit separately. To set SNMP get queries to each cluster unit you must create a special get command that includes the serial number of the cluster unit.

Alternatively you can use the HA reserved management interface feature to give each cluster unit a different management IP address. Then you can create an SNMP get command for each cluster unit that just includes the management IP address and does not have to include the serial number.

SNMP get command syntax for the primary unit

Normally, to get configuration and status information for a standalone FortiGate or for a primary unit, an SNMP manager would use an SNMP get commands to get the information in a MIB field. The SNMP get command syntax would be similar to the following:

```
snmpget -v2c -c <community_name> <address_ipv4> {<OID> | <MIB_field>}
```

where:

<community_name> is an SNMP community name added to the FortiGate configuration. You can add more than one community name to a FortiGate SNMP configuration. The most commonly used community name is public.

<address_ipv4> is the IP address of the FortiGate interface that the SNMP manager connects to.

{<OID> | <MIB_field>} is the object identifier (OID) for the MIB field or the MIB field name itself. The HA MIB fields and OIDs are listed below:

SNMP field names and OIDs

MIB field	OID	Description
fgHaSystemMode	.1.3.6.1.4.1.12356.101.13.1.1.0	HA mode (standalone, a-a, or a-p)
fgHaGroupId	.1.3.6.1.4.1.12356.101.13.1.2.0	The HA group ID of the cluster unit.
fgHaPriority	.1.3.6.1.4.1.12356.101.13.1.3.0	The HA priority of the cluster unit. Default 128.
fgHaOverride	.1.3.6.1.4.1.12356.101.13.1.4.0	Whether HA override is disabled or enabled for the cluster unit.
fgHaAutoSync	.1.3.6.1.4.1.12356.101.13.1.5.0	Whether automatic HA synchronization is disabled or enabled.

MIB field	OID	Description
fgHaSchedule	.1.3.6.1.4.1.12356.101.13.1.6.0	The HA load balancing schedule. Set to none unless operating in a-p mode.
fgHaGroupName	.1.3.6.1.4.1.12356.101.13.1.7.0	The HA group name.
fgHaStatsIndex	.1.3.6.1.4.1.12356.101.13.2.1.1.1.1	An index value that identifies the FortiGates in an HA cluster. The index is always 1 for the FortiGate that receives the HA get. The other FortiGate(s) in the cluster will have an index of 2, 3, or 4. For example, if you get the stats index from the primary FortiGate, the primary FortiGate will have a stats index of 1 and the backup FortiGate will have a stats index of 2. If you get the stats index from the backup unit, the backup unit will have a stats index of 1 and the primary unit will have a stats index of 2.
fgHaStatsSerial	.1.3.6.1.4.1.12356.101.13.2.1.1.2.1	The serial number of the cluster unit.
fgHaStatsCpuUsage	.1.3.6.1.4.1.12356.101.13.2.1.1.3.1	The cluster unit's current CPU usage.
fgHaStatsMemUsage	.1.3.6.1.4.1.12356.101.13.2.1.1.4.1	The cluster unit's current Memory usage.
fgHaStatsNetUsage	.1.3.6.1.4.1.12356.101.13.2.1.1.5.1	The cluster unit's current Network bandwidth usage.
fgHaStatsSesCount	.1.3.6.1.4.1.12356.101.13.2.1.1.6.1	The cluster unit's current session count.
fgHaStatsPktCount	.1.3.6.1.4.1.12356.101.13.2.1.1.7.1	The cluster unit's current packet count.
fgHaStatsByteCount	.1.3.6.1.4.1.12356.101.13.2.1.1.8.1	The cluster unit's current byte count.
fgHaStatsIdsCount	.1.3.6.1.4.1.12356.101.13.2.1.1.9.1	The number of attacks reported by the IPS for the cluster unit.
fgHaStatsAvCount	.1.3.6.1.4.1.12356.101.13.2.1.1.10.1	The number of viruses reported by the antivirus system for the cluster unit.
fgHaStatsHostname	.1.3.6.1.4.1.12356.101.13.2.1.1.11.1	The hostname of the cluster unit.

To get the HA priority for the primary unit

The following SNMP get command gets the HA priority for the primary unit. The community name is `public`. The IP address of the cluster interface configured for SNMP management access is `10.10.10.1`. The HA priority MIB field is `fgHaPriority` and the OID for this MIB field is `1.3.6.1.4.1.12356.101.13.1.3.0`. The first command uses the MIB field name and the second uses the OID:

```
snmpget -v2c -c public 10.10.10.1 fgHaPriority
```

```
snmpget -v2c -c public 10.10.10.1 1.3.6.1.4.1.12356.101.13.1.3.0
```

SNMP get command syntax for any cluster unit

To get configuration status information for a specific cluster unit (for the primary unit or for any subordinate unit), the SNMP manager must add the serial number of the cluster unit to the SNMP get command after the community name. The community name and the serial number are separated with a dash. The syntax for this SNMP get command would be:

```
snmpget -v2c -c <community_name>-<fgt_serial> <address_ipv4> {<OID> | <MIB_field>}
```

where:

<community_name> is an SNMP community name added to the FortiGate configuration. You can add more than one community name to a FortiGate SNMP configuration. All units in the cluster have the same community name. The most commonly used community name is `public`.

<fgt_serial> is the serial number of any cluster unit. For example, FGT4002803033172. You can specify the serial number of any cluster unit, including the primary unit, to get information for that unit.

<address_ipv4> is the IP address of the FortiGate interface that the SNMP manager connects to.

{<OID> | <MIB_field>} is the object identifier (OID) for the MIB field or the MIB field name itself.

If the serial number matches the serial number of a subordinate unit, the SNMP get request is sent over the HA heartbeat link to the subordinate unit. After processing the request, the subordinate unit sends the reply back over the HA heartbeat link back to the primary unit. The primary unit then forwards the response back to the SNMP manager.

If the serial number matches the serial number of the primary unit, the SNMP get request is processed by the primary unit. You can actually add a serial number to the community name of any SNMP get request. But normally you only need to do this for getting information from a subordinate unit.

To get the CPU usage for a subordinate unit

The following SNMP get command gets the CPU usage for a subordinate unit in a FortiGate-5001SX cluster. The subordinate unit has serial number FG50012205400050. The community name is `public`. The IP address of the FortiGate interface is 10.10.10.1. The HA status table MIB field is `fgHaStatsCpuUsage` and the OID for this MIB field is 1.3.6.1.4.1.12356.101.13.2.1.1.3.1. The first command uses the MIB field name and the second uses the OID for this table:

```
snmpget -v2c -c public-FG50012205400050 10.10.10.1 fgHaStatsCpuUsage
snmpget -v2c -c public-FG50012205400050 10.10.10.1 1.3.6.1.4.1.12356.101.13.2.1.1.3.1
```

FortiGate SNMP recognizes the community name with syntax `<community_name>-<fgt_serial>`. When the primary unit receives an SNMP get request that includes the community name followed by serial number, the FGCP extracts the serial number from the request. Then the primary unit redirects the SNMP get request to the cluster unit with that serial number. If the serial number matches the serial number of the primary unit, the SNMP get is processed by the primary unit.

Getting serial numbers of cluster units

The following SNMP get commands use the MIB field name `fgHaStatsSerial.<index>` to get the serial number of each cluster unit. Where `<index>` is the cluster unit's cluster index and 1 is the cluster index of the primary unit, 2 is the cluster index of the first subordinate unit, and 3 is the cluster index of the second subordinate unit.

The OID for this MIB field is 1.3.6.1.4.1.12356.101.13.2.1.1.2.1. The community name is `public`. The IP address of the FortiGate interface is 10.10.10.1.

The first command uses the MIB field name and the second uses the OID for this table and gets the serial number of the primary unit:

```
snmpget -v2c -c public 10.10.10.1 fgHaStatsSerial.1
snmpget -v2c -c public 10.10.10.1 1.3.6.1.4.1.12356.101.13.2.1.1.2.1
```

The second command uses the MIB field name and the second uses the OID for this table and gets the serial number of the first subordinate unit:

```
snmpget -v2c -c public 10.10.10.1 fgHaStatsSerial.2
snmpget -v2c -c public 10.10.10.1 1.3.6.1.4.1.12356.101.13.2.2.2
```

SNMP get command syntax - reserved management interface enabled

To get configuration and status information for any cluster unit where you have enabled the HA reserved management interface feature and assigned IP addresses to the management interface of each cluster unit, an SNMP manager would use the following get command syntax:

```
snmpget -v2c -c <community_name> <mgmt_address_ipv4> {<OID> | <MIB_field>}
```

where:

<community_name> is an SNMP community name added to the FortiGate configuration. You can add more than one community names to a FortiGate SNMP configuration. The most commonly used community name is public.

<mgmt_address_ipv4> is the IP address of the FortiGate HA reserved management interface that the SNMP manager connects to.

{<OID> | <MIB_field>} is the object identifier (OID) for the MIB field or the MIB field name itself. To find OIDs and MIB field names see your FortiGate's online help.

Adding FortiClient licenses to a cluster

Each FortiGate in a cluster must have its own FortiClient license. Contact your reseller to purchase FortiClient licenses for all of the FortiGates in your cluster.

When you receive the license keys you can visit the [Fortinet Support](#) website and add the FortiClient license keys to each FortiGate. Then, as long as the cluster can connect to the internet each cluster unit receives its FortiClient license key from the FortiGuard network.

Adding FortiClient licenses to cluster units with a reserved management interface

You can also use the following steps to manually add license keys to your cluster units from the GUI or CLI. Your cluster must be connected to the internet and you must have configured a reserved management interface for each cluster unit.

1. Log into the GUI of each cluster unit using its reserved management interface IP address.
2. Go to the **License Information** dashboard widget and beside **FortiClient** select **Enter License**.
3. Enter the license key and select **OK**.
4. Confirm that the license has been installed and the correct number of FortiClients are licensed.
5. Repeat for all of the cluster units.

You can also use the reserved management IP address to log into each cluster unit CLI and use following command to add the license key:

```
execute FortiClient-NAC update-registration-license <license-key>
```

You can connect to the CLIs of each cluster unit using their reserved management IP address.

Adding FortiClient licenses to cluster units with no reserved management interface

If you have not set up reserved management IP addresses for your cluster units, you can still add FortiClient license keys to each cluster unit. You must log into the primary unit and then use the `execute ha manage` command to connect to each cluster unit CLI. For example, use the following steps to add a FortiClient license key a cluster of three FortiGates:

1. Log into the primary unit CLI and enter the following command to confirm the serial number of the primary unit:

```
get system status
```

2. Add the FortiClient license key for that serial number to the primary unit:

```
execute FortiClient-NAC update-registration-license <license-key>
```

You can also use the GUI to add the license key to the primary unit.

3. Enter the following command to log into the first subordinate unit:

```
execute ha manage 1
```

4. Enter the following command to confirm the serial number of the cluster unit that you have logged into:

```
get system status
```

5. Add the FortiClient license key for that serial number to the cluster unit:

```
execute FortiClient-NAC update-registration-license <license-key>
```

6. Enter the following command to log into the second subordinate unit:

```
execute ha manage 2
```

7. Enter the following command to confirm the serial number of the cluster unit that you have logged into:

```
get system status
```

8. Add the FortiClient license key for that serial number to the cluster unit:

```
execute FortiClient-NAC update-registration-license <license-key>
```

Viewing FortiClient license status and active FortiClient users for each cluster unit

To view FortiClient license status and FortiClient information for each cluster unit you must log into each cluster unit's GUI or CLI. You can do this by connecting to each cluster unit's reserved management interface if they are configured. If you have not configured reserved management interfaces you can use the `execute ha manage` command to log into each cluster unit CLI.

From the GUI, view FortiClient License status from the License Information dashboard widget and select **Details** to display the list of active FortiClient users connecting through that cluster unit. You can also see active FortiClient users by going to **User & Device > Monitor > FortiClient**.

From the CLI you can use the `execute FortiClient {list | info}` command to display FortiClient license status and active FortiClient users.

For example, use the following command to display the FortiClient license status of the cluster unit that you are logged into:

```
execute forticlient info
Maximum FortiClient connections: unlimited.
Licensed connections: 114
  NAC: 114
  WANOPT: 0
  Test: 0
```

```
Other connections:
  IPsec: 0
  SSLVPN: 0
```

Use the following command to display the list of active FortiClient users connecting through the cluster unit. The output shows the time the connection was established, the type of FortiClient connection, the name of the device, the user name of the person connecting, the FortiClient ID, the host operating system, and the source IP address of the session.

```
execute forticlient list
TIMESTAMP TYPE CONNECT-NAME USER CLIENT-ID HOST-OS SRC-IP
20141017 09:13:33 NAC Gordon-PC Gordon 11F76E902611484A942E31439E428C5C Microsoft
  Windows 7 , 64-bit Service Pack 1 (build 7601) 172.20.120.10
20141017 09:11:55 NAC Gordon-PC 11F76E902611484A942E31439E428C5C Microsoft Windows 7 ,
  64-bit Service Pack 1 (build 7601) 172.20.120.10
20141017 07:27:11 NAC Desktop11 Richie 9451C0B8EE3740AEB7019E920BB3761B Microsoft
  Windows 7, 64-bit Service Pack 1 (build 7601) 172.20.120.20
```

Cluster members list

To display the cluster members list, go to **System > HA**.

The cluster members list displays illustrations of the front panels of the cluster units. If the network jack for an interface is shaded green, the interface is connected. Hover the mouse pointer over each illustration to view the cluster unit host name, serial number, and how long the unit has been operating (up time). The list of monitored interfaces is also displayed.

From the cluster members list you can:

- View HA statistics.
- Use the up and down arrows to change the order in which cluster units are listed.
- See the host name of each cluster unit. To change the primary unit host name, go to the system dashboard and select Change beside the current host name in the System Information widget. To view and change a subordinate unit host name, from the cluster members list select the edit icon for a subordinate unit.
- View the status or role of each cluster unit.
- View and optionally change the HA configuration of the operating cluster.
- View and optionally change the host name and device priority of a subordinate unit.
- Disconnect a cluster unit from a cluster.
- Download the Debug log for any cluster unit. You can send this debug log file to [Fortinet Support](#) to help diagnose problems with the cluster or with individual cluster units.

Virtual cluster members list

If virtual domains are enabled, you can display the cluster members list to view the status of the operating virtual clusters. The virtual cluster members list shows the status of both virtual clusters including the virtual domains added to each virtual cluster.

To display the virtual cluster members list for an operating cluster log in as the admin administrator, select Global Configuration and go to **System > HA**.

The functions of the virtual cluster members list are the same as the functions of the Cluster Members list with the following exceptions.

- When you select the edit icon for a primary unit in a virtual cluster, you can change the virtual cluster 1 and virtual cluster 2 device priority of this cluster unit and you can edit the VDOM partitioning configuration of the cluster.
- When you select the edit icon for a subordinate unit in a virtual cluster, you can change the device priority for the subordinate unit for the selected virtual cluster.

Also, the HA cluster members list changes depending on the cluster unit that you connect to.

Viewing HA statistics

From the cluster members list you can select **View HA statistics** to display the serial number, status, and monitor information for each cluster unit. To view HA statistics, go to **System > HA** and select View HA Statistics. Note the following about the HA statistics display:

- Use the serial number ID to identify each FortiGate in the cluster. The cluster ID matches the FortiGate serial number.
- Status indicates the status of each cluster unit. A green check mark indicates that the cluster unit is operating normally. A red X indicates that the cluster unit cannot communicate with the primary unit.
- The up time is the time in days, hours, minutes, and seconds since the cluster unit was last started.
- The GUI displays CPU usage for core processes only. CPU usage for management processes (for example, for HTTPS connections to the GUI) is excluded.
- The GUI displays memory usage for core processes only. Memory usage for management processes (for example, for HTTPS connections to the GUI) is excluded.

Changing the HA configuration of an operating cluster

To change the configuration settings of an operating cluster, go to **System > HA** to display the cluster members list. Select Edit for the master (or primary) unit in the cluster members list to display the HA configuration page for the cluster.

You can use the HA configuration page to check and fine tune the configuration of the cluster after the cluster is up and running. For example, if you connect or disconnect cluster interfaces you may want to change the Port Monitor configuration.

Any changes you make on this page, with the exception of changes to the device priority, are first made to the primary unit configuration and then synchronized to the subordinate units. Changing the device priority only affects the primary unit.

Changing the HA configuration of an operating virtual cluster

To change the configuration settings of the primary unit in a functioning cluster with virtual domains enabled, log in as the admin administrator, select Global Configuration and go to **System > HA** to display the cluster members list. Select Edit for the master (or primary) unit in virtual cluster 1 or virtual cluster 2 to display the HA configuration page for the virtual cluster.

You can use the virtual cluster HA configuration page to check and fine tune the configuration of both virtual clusters after the cluster is up and running. For example, you may want to change the Port Monitor configuration for virtual cluster 1 and virtual cluster 2 so that each virtual cluster monitors its own interfaces.

You can also use this configuration page to move virtual domains between virtual cluster 1 and virtual cluster 2. Usually you would distribute virtual domains between the two virtual clusters to balance the amount of traffic being processed by each virtual cluster.

Any changes you make on this page, with the exception of changes to the device priorities, are first made to the primary unit configuration and then synchronized to the subordinate unit.

You can also adjust device priorities to configure the role of this cluster unit in the virtual cluster. For example, to distribute traffic to both cluster units in the virtual cluster configuration, you would want one cluster unit to be the primary unit for virtual cluster 1 and the other cluster unit to be the primary unit for virtual cluster 2. You can create this configuration by setting the device priorities. The cluster unit with the highest device priority in virtual cluster 1 becomes the primary unit for virtual cluster 1. The cluster unit with the highest device priority in virtual cluster 2 becomes the primary unit in virtual cluster 2.

Changing the subordinate unit host name and device priority

To change the host name and device priority of a subordinate unit in an operating cluster, go to **System > HA** to display the cluster members list. Select Edit for any slave (subordinate) unit in the cluster members list.

To change the host name and device priority of a subordinate unit in an operating cluster with virtual domains enabled, log in as the admin administrator, select Global Configuration and go to **System > HA** to display the cluster members list. Select Edit for any slave (subordinate) unit in the cluster members list.

You can change the host name (Peer) and device priority (Priority) of this subordinate unit. These changes only affect the configuration of the subordinate unit.

The device priority is not synchronized among cluster members. In a functioning cluster you can change device priority to change the priority of any unit in the cluster. The next time the cluster negotiates, the cluster unit with the highest device priority becomes the primary unit.

The device priority range is 0 to 255. The default device priority is 128.

Upgrading cluster firmware

You can upgrade the FortiOS firmware running on an HA cluster in the same manner as upgrading the firmware running on a standalone FortiGate. During a normal firmware upgrade, the cluster upgrades the primary unit and all subordinate units to run the new firmware image. The firmware upgrade takes place without interrupting communication through the cluster.



Upgrading cluster firmware to a new major release (for example upgrading from 5.6.3 to 6.0.2) is supported for clusters. Make sure you are taking an appropriate [upgrade path](#). Even so you should back up your configuration and only perform such a firmware upgrade during a maintenance window.

To upgrade the firmware without interrupting communication through the cluster, the cluster goes through a series of steps that involve first upgrading the firmware running on the subordinate units, then making one of the subordinate units the primary unit, and finally upgrading the firmware on the former primary unit. These steps are transparent to the user and the network, but depending upon your HA configuration may result in the cluster selecting a new primary unit.

The following sequence describes in detail the steps the cluster goes through during a firmware upgrade and how different HA configuration settings may affect the outcome.

1. The administrator uploads a new firmware image from the GUI or CLI.
2. If the cluster is operating in active-active mode load balancing is turned off.
3. The cluster upgrades the firmware running on all of the subordinate units.
4. Once the subordinate units have been upgraded, a new primary unit is selected.

This primary unit will be running the new upgraded firmware.

5. The cluster now upgrades the firmware of the former primary unit.

If the age of the new primary unit is more than 300 seconds (5 minutes) greater than the age of all other cluster units, the new primary unit continues to operate as the primary unit.

This is the intended behavior but does not usually occur because the age difference of the cluster units is usually less than the cluster age difference margin of 300 seconds. So instead, the cluster negotiates again to select a primary unit as described in [Primary unit selection on page 1385](#).

You can keep the cluster from negotiating again by reducing the cluster age difference margin using the `ha-uptime-diff-margin` option. However, you should be cautious when reducing the age or other problems may occur. For information about the cluster age difference margin, see [Cluster age difference margin \(grace period\) on page 1388](#). For more information about changing the cluster age margin, see [Changing the cluster age difference margin on page 1388](#).

6. If the cluster is operating in active-active mode, load balancing is turned back on.



If, during the firmware upgrade process all of the subordinate units crash or otherwise stop responding, the primary unit will not be upgraded to the new firmware, but will continue to operate normally. The primary unit waits until at least one subordinate unit rejoins the cluster before upgrading its firmware.

Changing how the cluster processes firmware upgrades

By default cluster firmware upgrades proceed as uninterruptable upgrades that do not interrupt traffic flow. If required, you can use the following CLI command to change how the cluster handles firmware upgrades. You might want to change this setting if you are finding uninterruptable upgrades take too much time.

```
config system ha
    set uninterruptible-upgrade disable
end
```

`uninterruptible-upgrade` is enabled by default. If you disable `uninterruptible-upgrade` the cluster still upgrades the firmware on all cluster units, but all cluster units are upgraded at once; which takes less time but interrupts communication through the cluster.

Synchronizing the firmware build running on a new cluster unit

If the firmware build running on a FortiGate that you add to a cluster is older than the cluster firmware build, you may be able to use the following steps to synchronize the firmware running on the new cluster unit.

This procedure describes re-installing the same firmware build on a cluster to force the cluster to upgrade all cluster units to the same firmware build.

Due to firmware upgrade and synchronization issues, in some cases this procedure may not work. In all cases it will work to install the same firmware build on the new unit as the one that the cluster is running before adding the new unit to the cluster.

To synchronize the firmware build running on a new cluster unit

1. Obtain a firmware image that is the same as build already running on the cluster.
2. Connect to the cluster using the GUI.
3. Go to the **System Information** dashboard widget.
4. Select **Update** beside **Firmware Version**.

You can also install a newer firmware build.

5. Select **OK**.

After the firmware image is uploaded to the cluster, the primary unit upgrades all cluster units to this firmware build.

Downgrading cluster firmware

For various reasons you may need to downgrade the firmware that a cluster is running. You can use the information in this section to downgrade the firmware version running on a cluster.

In most cases you can downgrade the firmware on an operating cluster using the same steps as for a firmware upgrade. A warning message appears during the downgrade but the downgrade usually works and after the downgrade the cluster continues operating normally with the older firmware image.

Downgrading between some firmware versions, especially if features have changed between the two versions, may not always work without the requirement to fix configuration issues after the downgrade.

Only perform firmware downgrades during maintenance windows and make sure you back up your cluster configuration before the downgrade.

If the firmware downgrade that you are planning may not work without configuration loss or other problems, you can use the following downgrade procedure to make sure your configuration is not lost after the downgrade.

To downgrade cluster firmware

This example shows how to downgrade the cluster shown in Example NAT/Route mode HA network topology. The cluster consists of two cluster units (FGT_ha_1 and FGT_ha_2). The port1 and port2 interfaces are connected to networks and the port3 and port4 interfaces are connected together for the HA heartbeat.

This example, describes separating each unit from the cluster and downgrading the firmware for the standalone FortiGates. There are several ways you could disconnect units from the cluster. This example describes using the disconnect from cluster function on the cluster members list GUI page.

1. Go to the **System Information** dashboard widget and backup the cluster configuration.

From the CLI use `execute backup config`.

2. Go to **System > HA** and for FGT_ha_1 select the **Disconnect from cluster** icon.

3. Select the port2 interface and enter an IP address and netmask of 10.11.101.101/24 and select **OK**.

From the CLI you can enter the following command (FG600B3908600705 is the serial number of the cluster unit) to be able to manage the standalone FortiGate by connecting to the port2 interface with IP address and netmask 10.11.101.101/24.

```
execute ha disconnect FG600B3908600705 port2 10.11.101.101/24
```

After FGT_ha_1 is disconnected, FGT_ha_2 continues processing traffic.

4. Connect to the FGT_ha_1 GUI or CLI using IP address 10.11.101.101/24 and follow normal procedures to downgrade standalone FortiGate firmware.
5. When the downgrade is complete confirm that the configuration of 620_ha_1 is correct.
6. Set the HA mode of FGT_ha_2 to Standalone and follow normal procedures to downgrade standalone FortiGate firmware.

Network communication will be interrupted for a short time during the downgrade.

7. When the downgrade is complete confirm that the configuration of FGT_ha_2 is correct.
8. Set the HA mode of FGT_ha_2 to Active-Passive or the required HA mode.
9. Set the HA mode of FGT_ha_1 to the same mode as FGT_ha_2.
If you have not otherwise changed the HA settings of the cluster units and if the firmware downgrades have not affected the configurations the units should negotiate and form cluster running the downgraded firmware.

Backing up and restoring the cluster configuration

You can backup and restore the configuration of a cluster in the same way as backing up and restoring a standalone FortiGate unit. Backing up the cluster from the primary unit GUI or CLI saves a single configuration file for the cluster. If you restore this configuration file, the configuration of all cluster units is restored. The restore process keeps configuration settings of individual cluster units that are not synchronized unchanged but resets all other configuration setting to those in the restored configuration file.



When restoring the configuration of a cluster, all cluster units reboot to install the new configuration. This may result in a brief traffic interruption as all cluster units may restart at the same time.

Restoring settings that are not synchronized

The FGCP does not synchronize some FortiOS configuration settings. For details about settings that are not synchronized, see [Synchronizing the configuration \(and settings that are not synchronized\) on page 1376](#). If you need to restore the configuration of the cluster including the configuration settings that are not synchronized, you should first restore the configuration of the primary FortiGate and then restore the configuration of the other cluster units. Alternatively you could log into each FortiGate in the cluster and manually add the configuration settings that were not restored.

Monitoring cluster units for failover

If the primary unit in the cluster fails, the units in the cluster renegotiate to select a new primary unit. Failure of the primary unit results in the following:

- If SNMP is enabled, the new primary unit sends HA trap messages. The messages indicate a cluster status change, HA heartbeat failure, and HA member down.
- If event logging is enabled and HA activity event is selected, the new primary unit records log messages that show that the unit has become the primary unit.
- If alert email is configured to send email for HA activity events, the new primary unit sends an alert email containing the log message recorded by the event log.
- The cluster contains fewer FortiGates. The failed primary unit no longer appears on the Cluster Members list.
- The host name and serial number of the primary unit changes. You can see these changes when you log into the GUI or CLI.
- The cluster info displayed on the dashboard, cluster members list or from the `get system ha status` command changes.

If a subordinate unit fails, the cluster continues to function normally. Failure of a subordinate unit results in the following:

- If event logging is enabled and HA activity event is selected, the primary unit records log messages that show that a subordinate has been removed from the cluster.
- If alert email is configured to send email for HA activity events, the new primary unit sends an alert email containing the log message recorded by the event log.
- The cluster contains fewer FortiGates. The failed unit no longer appears on the Cluster Members list.

Viewing cluster status from the CLI

Use the `get system ha status` command to display information about an HA cluster. The command displays general HA configuration settings. The command also displays information about how the cluster unit that you have logged into is operating in the cluster. You can enter the `get system ha status` command from the primary or backup units. The output produced by the command is similar for each unit, it shows cluster data as well as data for the FortiGate that you are logged into.

For a virtual cluster configuration, the `get system ha status` command displays information about how the cluster unit that you have logged into is operating in virtual cluster 1 and virtual cluster 2. For example, if you connect to the cluster unit that is the primary unit for virtual cluster 1 and the subordinate unit for virtual cluster 2, the output of the `get system ha status` command shows virtual cluster 1 in the work state and virtual cluster 2 in the standby state. The `get system ha status` command also displays additional information about virtual cluster 1 and virtual cluster 2.

The command includes the following fields.

Field	Description
HA Health Status	Indicates if all cluster units are operating normally (OK) or if a problem was detected with the cluster. For example, a message similar to <code>ERROR <serial-number> is lost @ <date> <time></code> appears if one the subordinate units leaves the cluster.
Model	The FortiGate model number.
Mode	The HA mode of the cluster, for example, HA A-P or HA A-A.
Group	The group ID of the cluster.
Debug	The debug status of the cluster.
Cluster Uptime	The number of days, hours, minutes, and seconds that the cluster has been operating.
Cluster state changed time	The date and time at which the FortiGate most recently changed state. For example, the last time the FortiGate joined the cluster or changed from the primary unit to a backup unit, and so on.

Field	Description
Master selected using	Shows how the primary unit was selected the last four times that the cluster negotiated. For example, when a cluster first forms, this part of the command output could have one line showing that the primary unit is the cluster unit with the highest uptime. Up to four lines can be included as the cluster negotiates to choose a new primary unit on different occasions. Each line includes a time stamp and the criteria used to select the primary unit.
ses_pickup	The status of session pickup: enable or disable.
load_balance	The status of the <code>load-balance-all</code> keyword: enable or disable. Active-active clusters only.
load_balance_udp	The status of the <code>load-balance-udp</code> keyword: enable or disable. Available on some FortiGate models. Active-active clusters only.
schedule	The active-active load balancing schedule. Active-active clusters only.
override	The status of the override option for the current cluster unit: enable or disable.
Configuration Status	Shows if the configurations of each of the cluster units are synchronized or not.
System Usage stats	Shows how busy each cluster unit is by displaying the number of sessions being processed by the cluster unit, CPU usage, and memory usage.
HBDEV stats	Shows the status of each cluster unit's heartbeat interfaces. Includes whether the interfaces are up or down, how much data they have processed, as well as errors found.
Master Slave	<p>Displays the host name, serial number, and cluster index of the primary unit (master) and the subordinate units (slave). The FortiGate with cluster index 0 is the primary unit and the FortiGates with cluster indexes 1 to 3 are the backup units.</p> <p>The order in which the cluster units are listed starts with the cluster unit that you are logged into.</p>
number of vcluster	The number of virtual clusters. If virtual domains are not enabled, the cluster has one virtual cluster. If virtual domains are enabled the cluster has two virtual clusters.
vcluster 1 vcluster 2	The heartbeat interface IP address of the primary unit in each virtual cluster. If virtual domains are not enabled there is one vcluster and this is the IP address of the primary unit. If virtual domains are enabled then each vcluster line will have an IP address. If the IP addresses are the same then the same FortiGate is the primary unit for both virtual clusters.

Field	Description
<code>vcluster 1</code>	<p>The HA state (hello, work, or standby) and HA heartbeat IP address of the primary unit. If virtual domains are not enabled, <code>vcluster 1</code> displays information for the cluster. If virtual domains are enabled, <code>vcluster 1</code> displays information for virtual cluster 1.</p> <p><code>vcluster 1</code> also lists the primary unit (master) and subordinate units (slave) in virtual cluster 1. The list includes the serial number and operating cluster index of each cluster unit in virtual cluster 1. The cluster unit that you have logged into is at the top of the list. The FortiGate in the cluster with the highest serial number always has an operating cluster index of 0. Other FortiGates in the cluster get a higher operating cluster index based in their serial number. When you use the <code>execute ha manage</code> command to log into another FortiGate you use the operating cluster index to specify the FortiGate to log into.</p>
Master	<p>If virtual domains are not enabled and you connect to the primary unit CLI, the HA state of the cluster unit in virtual cluster 1 is work. The display lists the cluster units starting with the primary unit.</p>
Slave	<p>If virtual domains are not enabled and you connect to a subordinate unit CLI, the HA state of the cluster unit in virtual cluster 1 is standby. The display lists the cluster units starting with the subordinate unit that you have logged into.</p> <p>If virtual domains are enabled and you connect to the virtual cluster 1 primary unit CLI, the HA state of the cluster unit in virtual cluster 1 is work. The display lists the cluster units starting with the virtual cluster 1 primary unit.</p> <p>If virtual domains are enabled and you connect to the virtual cluster 1 subordinate unit CLI, the HA state of the cluster unit in virtual cluster 1 is standby. The display lists the cluster units starting with the subordinate unit that you are logged into.</p>
<code>vcluster 2</code>	<p><code>vcluster 2</code> only appears if virtual domains are enabled. <code>vcluster 2</code> displays the HA state (hello, work, or standby) and HA heartbeat IP address of the cluster unit that you have logged into in virtual cluster 2. The HA heartbeat IP address is 169.254.0.2 if you are logged into the primary unit of virtual cluster 2 and 169.254.0.1 if you are logged into a subordinate unit of virtual cluster 2.</p> <p><code>vcluster 2</code> also lists the primary unit (master) and subordinate units (slave) in virtual cluster 2. The list includes the cluster index and serial number of each cluster unit in virtual cluster 2. The cluster unit that you have logged into is at the top of the list.</p>
Master	<p>If you connect to the virtual cluster 2 primary unit CLI, the HA state of the cluster unit in virtual cluster 2 is <code>work</code>. The display lists the cluster units starting with the virtual cluster 2 primary unit.</p>
Slave	<p>If you connect to the virtual cluster 2 subordinate unit CLI, the HA state of the cluster unit in virtual cluster 2 is <code>standby</code>. The display lists the cluster units starting with the subordinate unit that you are logged into.</p>

Get system ha status example - two FortiGates in active-passive mode

The following example shows `get system ha status` output for a cluster of two FortiGate-600Ds operating in active-passive mode. The cluster is healthy and has been running for 3 hours and 26 minutes. Primary unit selection took place once and the cluster has been stable since then.

The following command output was produced by connecting to the primary unit CLI (host name Edge2-Primary).

```
get system ha status
HA Health Status: OK
Model: FortiGate-600D
Mode: HA A-P
Group: 25
Debug: 0
Cluster Uptime: 0 days 03:26:00
Cluster state change time: 2018-03-06 13:16:33
Master selected using:
    <2018/03/06 13:16:33> FGT6HD3916806098 is selected as the master because it has the
largest value of override priority.
    <2018/03/06 12:47:58> FGT6HD3916806070 is selected as the master because it has the
largest value of override priority.
    <2018/03/06 12:47:55> FGT6HD3916806098 is selected as the master because it has the
largest value of uptime.
    <2018/03/06 12:47:55> FGT6HD3916806098 is selected as the master because it's the only
member in the cluster.
ses_pickup: enable, ses_pickup_delay=disable
override: disable
Configuration Status:
    FGT6HD3916806098(updated 1 seconds ago): in-sync
    FGT6HD3916806070(updated 2 seconds ago): in-sync
System Usage stats:
    FGT6HD3916806098(updated 1 seconds ago):
        sessions=141, average-cpu-user/nice/system/idle=0%/0%/0%/100%, memory=34%
    FGT6HD3916806070(updated 2 seconds ago):
        sessions=12, average-cpu-user/nice/system/idle=0%/0%/0%/100%, memory=33%
HBDEV stats:
    FGT6HD3916806098(updated 1 seconds ago):
        port3: physical/1000full, up, rx-bytes/packets/dropped/errors=45437370/71531/0/0,
tx=36186194/65035/0/0
        port4: physical/1000full, up, rx-bytes/packets/dropped/errors=27843923/39221/0/0,
tx=27510707/39075/0/0
    FGT6HD3916806070(updated 2 seconds ago):
        port3: physical/1000full, up, rx-bytes/packets/dropped/errors=37267057/67136/0/0,
tx=46354380/73516/0/0
        port4: physical/1000full, up, rx-bytes/packets/dropped/errors=28294029/40177/0/0,
tx=28536766/40208/0/0
Master: Edge2-Primary , FGT6HD3916806098, cluster index = 0
Slave : Edge2-Backup , FGT6HD3916806070, cluster index = 1
number of vcluster: 1
vcluster 1: work 169.254.0.1
Master: FGT6HD3916806098, operating cluster index = 0
Slave : FGT6HD3916806070, operating cluster index = 1
```

The following command output was produced by using `execute ha manage 1` to log into the subordinate unit CLI of the cluster shown in the previous example. The host name of the subordinate unit is Edge2-Backup.

```

get system ha status
HA Health Status: OK
Model: FortiGate-600D
Mode: HA A-P
Group: 25
Debug: 0
Cluster Uptime: 0 days 03:33:04
Cluster state change time: 2018-03-06 13:16:33
Master selected using:
    <2018/03/06 13:16:33> FGT6HD3916806098 is selected as the master because it
has the largest value of override priority.
    <2018/03/06 12:47:58> FGT6HD3916806070 is selected as the master because it
has the largest value of override priority.
    <2018/03/06 12:47:57> FGT6HD3916806098 is selected as the master because it
has the largest value of uptime.
    <2018/03/06 12:47:56> FGT6HD3916806098 is selected as the master because it
has the largest value of uptime.
ses_pickup: enable, ses_pickup_delay=disable
override: disable
Configuration Status:
    FGT6HD3916806070 (updated 1 seconds ago): in-sync
    FGT6HD3916806098 (updated 1 seconds ago): in-sync
System Usage stats:
    FGT6HD3916806070 (updated 1 seconds ago):
        sessions=20, average-cpu-user/nice/system/idle=0%/0%/0%/100%, memory=34%
    FGT6HD3916806098 (updated 1 seconds ago):
        sessions=163, average-cpu-user/nice/system/idle=0%/0%/0%/100%, memory=34%
HBDEV stats:
    FGT6HD3916806070 (updated 1 seconds ago):
        port3: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=40755112/71809/0/0, tx=48104698/76943/0/0
        port4: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=29804904/42302/0/0, tx=30030641/42333/0/0
    FGT6HD3916806098 (updated 1 seconds ago):
        port3: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=47188898/74965/0/0, tx=39680065/69723/0/0
        port4: physical/1000full, up, rx-bytes/pack-
ets/dropped/errors=29338501/41347/0/0, tx=29022293/41201/0/0
Slave : Edge2-Backup , FGT6HD3916806070, cluster index = 1
Master: Edge2-Primary , FGT6HD3916806098, cluster index = 0
number of vcluster: 1
vcluster 1: standby 169.254.0.1
Slave : FGT6HD3916806070, operating cluster index = 1
Master: FGT6HD3916806098, operating cluster index = 0

```

About the HA operating cluster index and the execute ha manage command

When a cluster starts up, if primary unit select is based on serial number, the FortiGate Cluster Protocol (FGCP) assigns a cluster index and an HA heartbeat IP address to each cluster unit based on the serial number of the cluster unit:

- The FGCP selects the cluster unit with the highest serial number to become the primary unit. The FGCP assigns a cluster index of 0, an operating cluster index of 0, and an HA heartbeat IP address of 169.254.0.1 to this unit.

- The FGCP assigns a cluster index of 1, an operating cluster index of 1, and an HA heartbeat IP address of 169.254.0.2 to the cluster unit with the second highest serial number.
- If the cluster contains more units, the cluster unit with the third highest serial number is assigned a cluster index of 2, and operating cluster index of 2, and an HA heartbeat IP address of 169.254.0.3, and so on.

You can display the cluster index and operating cluster index assigned to each cluster unit using the `get system ha status` command. When you use the `execute ha manage` command you select a cluster unit to log into by entering its operating cluster index.

The operating cluster index and HA heartbeat IP address only change if a unit leaves the cluster or if a new unit joins the cluster. When one of these events happens, the FGCP resets the cluster index, operating cluster index, and HA heartbeat IP address of each cluster unit according to serial number in the same way as when the cluster first starts up.

If FortiGates don't leave or join, each cluster unit keeps its assigned operating cluster index, and HA heartbeat IP address since these are based on the FortiGate serial number, even as the units take on different roles in the cluster. After the operating cluster index and HA heartbeat IP addresses are set according to serial number, the FGCP checks other primary unit selection criteria such as device priority and monitored interfaces. Checking these criteria could result in selecting a cluster unit without the highest serial number to operate as the primary unit.

Even if the cluster unit without the highest serial number now becomes the primary unit, the operating cluster indexes and HA heartbeat IP addresses assigned to the individual cluster units do not change. Instead the FGCP changes the cluster index to reflect this role change. The cluster index is always 0 for the primary unit and 1 and higher for the other units in the cluster. By default both sets of cluster indexes are the same. But if primary unit selection selects the cluster unit that does not have the highest serial number to be the primary unit, then this cluster unit is assigned a cluster index of 0.

Using the `execute ha manage` command

When you use the CLI command `execute ha manage <index_integer>` to connect to the CLI of another cluster unit, the `<index_integer>` that you enter is the operating cluster index of the unit that you want to connect to.

Using `get system ha status` to display cluster indexes

You can display the cluster index assigned to each cluster unit using the CLI command `get system ha status`. The following example shows the information displayed by the `get system ha status` command for a cluster consisting of two FortiGates operating in active-passive HA mode with virtual domains not enabled and without virtual clustering.

```
get system ha status
.
.
.
Master: Edge2-Primary , FGT6HD3916806098, cluster index = 0
Slave : Edge2-Backup , FGT6HD3916806070, cluster index = 1
number of vcluster: 1
vcluster 1: work 169.254.0.1
Master: FGT6HD3916806098, operating cluster index = 0
Slave : FGT6HD3916806070, operating cluster index = 1
```

In this example, the cluster unit with serial number FGT6HD3916806098 has the highest serial number and so has a cluster index and an operating cluster index of 0 and the cluster unit with serial number

FGT6HD3916806070 has a cluster index and an operating cluster index of 1. From the CLI of the primary unit of this cluster you can connect to the CLI of the subordinate unit using the following command:

```
execute ha manage 1
```

This works because the cluster unit with serial number FGT6HD3916806070 has a cluster index of 1.

The last three lines of the command output display the status of vcluster 1. In a cluster consisting of two cluster units operating without virtual domains enabled, all clustering actually takes place in virtual cluster 1. HA is designed to work this way to support virtual clustering. If this cluster was operating with virtual domains enabled, adding virtual cluster 2 is similar to adding a new copy of virtual cluster 1. Virtual cluster 2 is visible in the `get system ha status` command output when you add virtual domains to virtual cluster 2.

The HA heartbeat IP address displayed by the command is the HA heartbeat IP address of the cluster unit that is actually operating as the primary unit. For a default configuration, this IP address will always be 169.254.0.1 because the cluster unit with the highest serial number will be the primary unit. This IP address changes if the operating primary unit is not the primary unit with the highest serial number.

Example where the cluster index and operating cluster index do not match

This example shows `get system ha status` command output for the same cluster. However, in this example the device priority of the cluster unit with the serial number FGT6HD3916806098 is increased to 250. As a result the cluster unit with the lowest serial number becomes the primary unit. This means the cluster index and the operating cluster index of the cluster units do not match.

```
get system ha status
.
.
.
Master: Edge2-Primary , FGT6HD3916806098, cluster index = 1
Slave : Edge2-Backup , FGT6HD3916806070, cluster index = 0
number of vcluster: 1
vcluster 1: work 169.254.0.2
Master: FGT6HD3916806098, operating cluster index = 0
Slave : FGT6HD3916806070, operating cluster index = 1
```

The actual cluster indexes have not changed but the operating cluster indexes have. Also, the HA heartbeat IP address displayed for vcluster 1 has changed to 169.254.0.2.

Virtual clustering example output

The `get system ha status` command output is the same if a cluster is operating with virtual clustering turned on but with all virtual domains in virtual cluster 1. The following `get system ha status` command output example shows the same cluster operating as a virtual cluster with virtual domains in virtual cluster 1 and added to virtual cluster 2. In this example the cluster unit with serial number FG50012204400045 is the primary unit for virtual cluster 1 and the cluster unit with serial number FG50012205400050 is the primary unit for virtual cluster 2.

```
get system ha status
.
.
.
number of vcluster: 2
vcluster 1: work 169.254.0.2
Master: FG50012205400050, operating cluster index = 1
Slave : FG50012204400045, operating cluster index = 0
vcluster 2: standby 169.254.0.1
```

```
Master: FG50012205400050, operating cluster index = 0
```

```
Slave : FG50012204400045, operating cluster index = 1
```

This example shows three sets of indexes. The indexes in lines six and seven are still used by the `execute ha manage` command. The indexes on lines ten and eleven are for the primary and subordinate units in virtual cluster 1 and the indexes on the last two lines are for virtual cluster 2.

Managing individual cluster units

The following procedure describes how to use SSH to log into the primary unit CLI and from there to use the `execute ha manage` command to connect to the CLI of any other unit in the cluster. The procedure is very similar if you use telnet, or the GUI dashboard CLI console.

You can use the `execute ha manage` command from the CLI of any cluster unit to log into the CLI of another the cluster unit. Usually you would use this command from the CLI of the primary unit to log into the CLI of a subordinate unit. However, if you have logged into a subordinate unit CLI, you can use this command to log into the primary unit CLI, or the CLI of another subordinate unit.

Using SSH or telnet or the GUI CLI console you can only log into the primary unit CLI. Using a direct console connection you can log into any cluster unit. In both cases you can use `execute ha manage` to connect to the CLI of other cluster units.

1. Log into the primary unit CLI.

Connect to any cluster interface configured for SSH administrative access to log into the cluster.

2. Enter the following command followed by a space and type a question mark (?):

```
execute ha manage
```

The CLI displays a list of the serial numbers of all of the subordinate units in the cluster. Each cluster unit is numbered. The number is the operating cluster index.

3. Complete the command with the operating cluster index number of the subordinate unit to log into. For example, to log into subordinate unit 1, enter the following command:

```
execute ha manage 1
```

4. Log into the CLI of the selected subordinate unit.

The CLI prompt changes to the host name of the subordinate unit. You can use CLI commands to manage this subordinate unit. If you make changes to the configuration of any cluster unit (primary or subordinate unit) these changes are synchronized to all cluster units.

5. You can now use the `execute ha manage` command to connect to any other cluster unit (including the primary unit). You can also use the `exit` command to return to the primary unit CLI.

Disconnecting a cluster unit from a cluster

Use the following procedures to disconnect a cluster unit from a functioning cluster without disrupting the operation of the cluster. You can disconnect a cluster unit if you need to use the disconnected FortiGate for another purpose, such as to act as a standalone firewall.

You can use the following procedures for a standard cluster and for a virtual clustering configuration. To use the following procedures from a virtual cluster you must be logged in as the admin administrator and you must have selected Global Configuration.

When you disconnect a cluster unit you must assign an IP address and netmask to one of the interfaces of the disconnected unit. You can disconnect any unit from the cluster even the primary unit. After the unit is

disconnected, the cluster responds as if the disconnected unit has failed. The cluster may renegotiate and may select a new primary unit.

When the cluster unit is disconnected the HA mode is changed to standalone. In addition, all interface IP addresses of the disconnected unit are set to 0.0.0.0 except for the interface that you configure.

Otherwise the configuration of the disconnected unit is not changed. The HA configuration of the disconnected unit is not changed either (except to change the HA mode to Standalone).

To disconnect a cluster unit from a cluster - GUI

1. Go to **System > HA** to view the cluster members list.
2. Select the Disconnect from cluster icon for the cluster unit to disconnect from the cluster.
3. Select the interface that you want to configure. You also specify the IP address and netmask for this interface. When the FortiGate is disconnected, all management access options are enabled for this interface.
4. Specify an IP address and netmask for the interface. You can use this IP address to connect to the interface to configure the disconnected FortiGate.
5. Select **OK**.

The FortiGate is disconnected from the cluster and the cluster may renegotiate and select a new primary unit. The selected interface of the disconnected unit is configured with the specified IP address and netmask.

To disconnect a cluster unit from a cluster - CLI

1. Enter the following command to disconnect a cluster unit with serial number FGT5002803033050. The internal interface of the disconnected unit is set to IP address 1.1.1.1 and netmask 255.255.255.0

```
execute ha disconnect FGT5002803033050 internal 1.1.1.1 255.255.255.0
```

Adding a disconnected FortiGate back to its cluster

If you disconnect a FortiGate from a cluster, you can re-connect the disconnected FortiGate to the cluster by setting the HA mode of the disconnected unit to match the HA mode of the cluster. Usually the disconnected unit rejoins the cluster as a subordinate unit and the cluster automatically synchronizes its configuration.



You do not have to change the HA password on the disconnected unit unless the HA password has been changed after the unit was disconnected. Disconnecting a unit from a cluster does not change the HA password.



You should make sure that the device priority of the disconnected unit is lower than the device priority of the current primary unit. You should also make sure that the HA `override` CLI option is not enabled on the disconnected unit. Otherwise, when the disconnected unit joins the cluster, the cluster will renegotiate and the disconnected unit may become the primary unit. If this happens, the configuration of the disconnected unit is synchronized to all other cluster units. This configuration change might disrupt the operation of the cluster.

The following procedure assumes that the disconnected FortiGate is correctly physically connected to your network and to the cluster but is not running in HA mode and not part of the cluster.

Before you start this procedure you should note the device priority of the primary unit.

To add a disconnected FortiGate back to its cluster - GUI

1. Log into the disconnected FortiGate.
If virtual domains are enabled, log in as the admin administrator and select Global Configuration.
2. Go to **System > HA**.
3. Change Mode to match the mode of the cluster.
4. If required, change the group name and password to match the cluster.
5. Set the Device Priority lower than the device priority of the primary unit.
6. Select **OK**.
The disconnected FortiGate joins the cluster.

To add a disconnected FortiGate back to its cluster - CLI

1. Log into the CLI of the FortiGate to be added back to the cluster.
2. Enter the following command to access the global configuration and add the FortiGate back to a cluster operating in active-passive mode and set the device priority to 50 (a low number) so that this unit will not become the primary unit:

```
config global
  config system ha
    set mode a-p
    set priority 50
  end
end
```

You may have to also change the group name, group id and password. However if you have not changed these for the cluster or the FortiGate after it was disconnected from the cluster you should not have to adjust them now.

diagnose sys ha dump-by command

You can use the following diagnose command to display data about a cluster:

```
diagnose sys ha dump-by {group | vcluster | rcache | debug-zone | vdom | kernel | device}
```

kernel

This command displays the HA configuration stored by the kernel.

```
diagnose sys ha dump-by kernel
      HA information.
group_id=88, nvcluster=2, mode=2, load_balance=0, schedule=3, ldb_udp=0.
nvcluster=2, mode=2, ses_pickup=0, delay=0, load_balance=0
schedule=3, ldb_udp=0, standalone_ha=0, upgrade_mode=0.
vcluster 1:
FGT51E5618000206, 0, 0.
FGT51E5618000259, 1, 1.
vcluster 2:
FGT51E5618000206, 1, 1.
FGT51E5618000259, 0, 0.
```

HA and failover protection

In FortiGate active-passive HA, the FortiGate Clustering Protocol (FGCP) provides failover protection. This means that an active-passive cluster can provide FortiGate services even when one of the cluster units encounters a problem that would result in complete loss of connectivity for a stand-alone FortiGate. This failover protection provides a backup mechanism that can be used to reduce the risk of unexpected downtime, especially in a mission-critical environment.

The FGCP supports three kinds of failover protection. Device failover automatically replaces a failed device and restarts traffic flow with minimal impact on the network. Link failover maintains traffic flow if a link fails. Session failover resumes communication sessions with minimal loss of data if a device or link failover occurs.

This chapter describes how FGCP failover protection works and provides detailed NAT/Route and transparent mode packet flow descriptions.

About active-passive failover

To achieve failover protection in an active-passive cluster, one of the cluster units functions as the primary unit, while the rest of the cluster units are subordinate units, operating in an active stand-by mode. The cluster IP addresses and HA virtual MAC addresses are associated with the cluster interfaces of the primary unit. All traffic directed at the cluster is actually sent to and processed by the primary unit.

While the cluster is functioning, the primary unit functions as the FortiGate network security device for the networks that it is connected to. In addition, the primary unit and subordinate units use the HA heartbeat to keep in constant communication. The subordinate units report their status to the cluster unit and receive and store connection and state table updates.

Device failure

If the primary unit encounters a problem that is severe enough to cause it to fail, the remaining cluster units negotiate to select a new primary unit. This occurs because all of the subordinate units are constantly waiting to negotiate to become primary units. Only the heartbeat packets sent by the primary unit keep the subordinate units from becoming primary units. Each received heartbeat packet resets negotiation timers in the subordinate units. If this timer is allowed to run out because the subordinate units do not receive heartbeat packets from the primary unit, the subordinate units assume that the primary unit has failed, and negotiate to become primary units themselves.

Using the same FGCP negotiation process that occurs when the cluster starts up, after they determine that the primary unit has failed, the subordinate units negotiate amongst themselves to select a new primary unit. The subordinate unit that wins the negotiation becomes the new primary unit with the same MAC and IP addresses as the former primary unit. The new primary unit then sends gratuitous ARP packets out all of its interfaces to inform attached switches to send traffic to the new primary unit. Sessions then resume with the new primary unit.

Link failure

If a primary unit interface fails or is disconnected while a cluster is operation, a link failure occurs. When a link failure occurs the cluster units negotiate to select a new primary unit. Since the primary unit has not stopped operating, it participates in the negotiation. The link failure means that a new primary unit must be selected and the cluster unit with the link failure joins the cluster as a subordinate unit.

Just as for a device failover, the new primary unit sends gratuitous arp packets out all of its interfaces to inform attached switches to send traffic to it. Sessions then resume with the new primary unit.

If a subordinate unit experiences a device failure its status in the cluster does not change. However, in future negotiations a cluster unit with a link failure is unlikely to become the primary unit.

Session failover

If you enable session failover (also called session pickup) for the cluster, during cluster operation the primary unit informs the subordinate units of changes to the primary unit connection and state tables, keeping the subordinate units up-to-date with the traffic currently being processed by the cluster.

After a failover the new primary unit recognizes open sessions that were being handled by the cluster. The sessions continue to be processed by the new primary unit and are handled according to their last known state.

If you leave session pickup disabled, the cluster does not keep track of sessions and after a failover, active sessions have to be restarted or resumed.

Primary unit recovery

If a primary unit recovers after a device or link failure, it will operate as a subordinate unit, unless the `override` CLI keyword is enabled and its device priority is set higher than the unit priority of other cluster units (see [HA override on page 1393](#)).

About active-active failover

HA failover in a cluster running in active-active mode is similar to active-passive failover described above. Active-active subordinate units are constantly waiting to negotiate to become primary units and, if session failover is enabled, continuously receive connection state information from the primary unit. If the primary unit fails, or one of the primary unit interfaces fails, the cluster units use the same mechanisms to detect the failure and to negotiate to select a new primary unit. If session failover is enabled, the new primary unit also maintains communication sessions through the cluster using the shared connection state table.

Active-active HA load balances sessions among all cluster units. For session failover, the cluster must maintain all of these sessions. To load balance sessions, the functioning cluster uses a load balancing schedule to distribute sessions to all cluster units. The shared connection state table tracks the communication sessions being processed by all cluster units (not just the primary unit). After a failover, the new primary unit uses the load balancing schedule to re-distribute all of the communication sessions recorded in the shared connection state table among all of the remaining cluster units. The connections continue to be processed by the cluster, but possibly by a different cluster unit, and are handled according to their last known state.

Device failover

The FGCP provides transparent device failover. Device failover is a basic requirement of any highly available system. Device failover means that if a device fails, a replacement device automatically takes the place of the failed device and continues operating in the same manner as the failed device.

In the case of FortiOS HA, the device is the primary unit. If the primary unit fails, device failover ensures that one of the subordinate units in the cluster automatically takes the place of the primary unit and can continue processing network traffic in the same way as the failed primary unit.



Device failover does not maintain communication sessions. After a device failover, communication sessions have to be restarted. To maintain communication sessions, you must enable session failover. See [Device failover on page 1522](#).

FortiGate HA device failover is supported by the HA heartbeat, virtual MAC addresses, configuration synchronization, route synchronization and IPsec VPN SA synchronization.

The HA heartbeat makes sure that the subordinate units detect a primary unit failure. If the primary unit fails to respond on time to HA heartbeat packets the subordinate units assume that the primary unit has failed and negotiate to select a new primary unit.

The new primary unit takes the place of the failed primary unit and continues functioning in the same way as the failed primary unit. For the new primary unit to continue functioning like the failed primary unit, the new primary unit must be able to reconnect to network devices and the new primary unit must have the same configuration as the failed primary unit.

FortiGate HA uses virtual MAC addresses to reconnect the new primary unit to network devices. The FGCP causes the new primary unit interfaces to acquire the same virtual MAC addresses as the failed primary unit. As a result, the new primary unit has the same network identity as the failed primary unit.

The new primary unit interfaces have different physical connections than the failed primary unit. Both the failed and the new primary unit interfaces are connected to the same switches, but the new primary unit interfaces are connected to different ports on these switches. To make sure that the switches send packets to the new primary unit, the new primary unit interfaces send gratuitous ARP packets to the connected switches. These gratuitous ARP packets notify the switches that the primary unit MAC and IP addresses are on different switch ports and cause the switches to send packets to the ports connected to the new primary unit. In this way, the new primary unit continues to receive packets that would otherwise have been sent to the failed primary unit.

Configuration synchronization means that the new primary unit always has the same configuration as the failed primary unit. As a result the new primary unit operates in exactly the same way as the failed primary unit. If configuration synchronization were not available the new primary unit may not process network traffic in the same way as the failed primary unit.

Kernel routing table synchronization synchronizes the primary unit kernel routing table to all subordinate units so that after a failover the new primary unit does not have to form a completely new routing table. IPsec VPN SA synchronization synchronizes IPsec VPN security associations (SAs) and other IPsec session data so that after a failover the new primary unit can resume IPsec tunnels without having to establish new SAs.

HA heartbeat and communication between cluster units

The HA heartbeat keeps cluster units communicating with each other. The heartbeat consists of hello packets that are sent at regular intervals by the heartbeat interface of all cluster units. These hello packets describe the state of the cluster unit and are used by other cluster units to keep all cluster units synchronized.

HA heartbeat packets are non-TCP packets that use Ethertype values 0x8890, 0x8891, and 0x8893. The default time interval between HA heartbeats is 200 ms. The FGCP uses link-local IPv4 addresses in the 169.254.0.x range for HA heartbeat interface IP addresses.

For best results, isolate the heartbeat devices from your user networks by connecting the heartbeat devices to a separate switch that is not connected to any network. If the cluster consists of two FortiGates you can connect the heartbeat device interfaces directly using a crossover cable. Heartbeat packets contain sensitive information

about the cluster configuration. Heartbeat packets may also use a considerable amount of network bandwidth. For these reasons, it is preferable to isolate heartbeat packets from your user networks.

On startup, a FortiGate configured for HA operation broadcasts HA heartbeat hello packets from its HA heartbeat interface to find other FortiGates configured to operate in HA mode. If two or more FortiGates operating in HA mode connect with each other, they compare HA configurations (HA mode, HA password, and HA group ID). If the HA configurations match, the units negotiate to form a cluster.

While the cluster is operating, the HA heartbeat confirms that all cluster units are functioning normally. The heartbeat also reports the state of all cluster units, including the communication sessions that they are processing.

Heartbeat interfaces

A heartbeat interface is an Ethernet network interface in a cluster that is used by the FGCP for HA heartbeat communications between cluster units.

To change the HA heartbeat configuration go to **System > HA** and select the *FortiGate interfaces* to use as HA heartbeat interfaces.



Do not use a switch port for the HA heartbeat traffic. This configuration is not supported.

From the CLI enter the following command to make port4 and port5 HA heartbeat interfaces and give both interfaces a heartbeat priority of 150:

```
config system ha
    set hbdev port4 150 port5 150
end
```

The following example shows how to change the default heartbeat interface configuration so that the port4 and port1 interfaces can be used for HA heartbeat communication and to give the port4 interface the highest heartbeat priority so that port4 is the preferred HA heartbeat interface.

```
config system ha
    set hbdev port4 100 port1 50
end
```

By default, for most FortiGate models two interfaces are configured to be heartbeat interfaces. You can change the heartbeat interface configuration as required. For example you can select additional or different heartbeat interfaces. You can also select only one heartbeat interface.

In addition to selecting the heartbeat interfaces, you also set the **Priority** for each heartbeat interface. In all cases, the heartbeat interface with the highest priority is used for all HA heartbeat communication. If the interface fails or becomes disconnected, the selected heartbeat interface that has the next highest priority handles all heartbeat communication.

If more than one heartbeat interface has the same priority, the heartbeat interface with the highest priority that is also highest in the heartbeat interface list is used for all HA heartbeat communication. If this interface fails or becomes disconnected, the selected heartbeat interface with the highest priority that is next highest in the list handles all heartbeat communication.

The default heartbeat interface configuration sets the priority of two heartbeat interfaces to 50. You can accept the default heartbeat interface configuration if one or both of the default heartbeat interfaces are connected. You can select different heartbeat interfaces, select more heartbeat interfaces and change heartbeat priorities according to your requirements.

For the HA cluster to function correctly, you must select at least one heartbeat interface and this interface of all of the cluster units must be connected together. If heartbeat communication is interrupted and cannot failover to a second heartbeat interface, the cluster units will not be able to communicate with each other and more than one cluster unit may become a primary unit. As a result the cluster stops functioning normally because multiple devices on the network may be operating as primary units with the same IP and MAC addresses creating a kind of split brain scenario.

The heartbeat interface priority range is 0 to 512. The default priority when you select a new heartbeat interface is 0. The higher the number the higher the priority.

In most cases you can maintain the default heartbeat interface configuration as long as you can connect the heartbeat interfaces together. Configuring HA heartbeat interfaces is the same for virtual clustering and for standard HA clustering.

You can enable heartbeat communications for physical interfaces, but not for VLAN subinterfaces, IPsec VPN interfaces, redundant interfaces, or for 802.3ad aggregate interfaces. You cannot select these types of interfaces in the heartbeat interface list.

Selecting more heartbeat interfaces increases reliability. If a heartbeat interface fails or is disconnected, the HA heartbeat fails over to the next heartbeat interface.

You can select up to 8 heartbeat interfaces. This limit only applies to FortiGates with more than 8 physical interfaces.

HA heartbeat traffic can use a considerable amount of network bandwidth. If possible, enable HA heartbeat traffic on interfaces used only for HA heartbeat traffic or on interfaces connected to less busy networks.

Connecting HA heartbeat interfaces

For most FortiGate models if you do not change the heartbeat interface configuration, you can isolate the default heartbeat interfaces of all of the cluster units by connecting them all to the same switch. Use one switch per heartbeat interface. If the cluster consists of two units you can connect the heartbeat interfaces together using crossover cables.

HA heartbeat and data traffic are supported on the same cluster interface. In NAT/Route mode, if you decide to use heartbeat interfaces for processing network traffic or for a management connection, you can assign the interface any IP address. This IP address does not affect HA heartbeat traffic.

In transparent mode, you can connect the heartbeat interface to your network and enable management access. You would then establish a management connection to the interface using the transparent mode management IP address. This configuration does not affect HA heartbeat traffic.

Heartbeat packets and heartbeat interface selection

HA heartbeat hello packets are constantly sent by all of the enabled heartbeat interfaces. Using these hello packets, each cluster unit confirms that the other cluster units are still operating. The FGCP selects one of the heartbeat interfaces to be used for communication between the cluster units. The FGCP selects the heartbeat interface for heartbeat communication based on the linkfail states of the heartbeat interfaces, on the priority of the heartbeat interfaces, and on the interface index.

The FGCP checks the linkfail state of all heartbeat interfaces to determine which ones are connected. The FGCP selects one of these connected heartbeat interfaces to be the one used for heartbeat communication. The FGCP selects the connected heartbeat interface with the highest priority for heartbeat communication.

If more than one connected heartbeat interface has the highest priority the FGCP selects the heartbeat interface with the lowest interface index. The GUI lists the FortiGate interfaces in alphabetical order. This order corresponds to the interface index order with lowest index at the top and highest at the bottom. If more than one heartbeat interface has the highest priority, the FGCP selects the interface that is highest in the heartbeat interface list (or first in alphabetical order) for heartbeat communication.

If the interface that is processing heartbeat traffic fails or becomes disconnected, the FGCP uses the same criteria to select another heartbeat interface for heartbeat communication. If the original heartbeat interface is fixed or reconnected, the FGCP again selects this interface for heartbeat communication.

The HA heartbeat communicates cluster session information, synchronizes the cluster configuration, synchronizes the cluster kernel routing table, and reports individual cluster member status. The HA heartbeat constantly communicates HA status information to make sure that the cluster is operating properly.

Interface index and display order

The GUI and CLI display interface names in alphanumeric order. For example, the sort order for a FortiGate with 10 interfaces (named port1 through port10) places port10 at the bottom of the list:

- port1
- port2 through 9
- port10

However, interfaces are indexed in hash map order, rather than purely by alphabetic order or purely by interface number value comparisons. As a result, the list is sorted primarily alphabetical by interface name (for example, base1 is before port1), then secondarily by index numbers:

- port1
- port10
- port2 through port9

HA heartbeat interface IP addresses

The FGCP uses link-local IPv4 addresses ([RFC 3927](#)) in the 169.254.0.x range for HA heartbeat interface IP addresses and for inter-VDOM link interface IP addresses. When a cluster initially starts up, the primary unit heartbeat interface IP address is 169.254.0.1. Subordinate units are assigned heartbeat interface IP addresses in the range 169.254.0.2 to 169.254.0.63. HA inter-VDOM link interfaces on the primary unit are assigned IP addresses 169.254.0.65 and 169.254.0.66.

If a failover occurs, the primary unit heartbeat interface could be something other than 169.254.0.1. If for example, the first subordinate unit is now the primary unit, the primary unit heartbeat interface IP address would be 169.254.0.2.

The output from the `get system ha status` CLI command shows the HA heartbeat interface IP address of the primary unit.

```
get system ha status
.
.
.
vcluster 1: work 169.254.0.2
.
.
.
```

You can also use the `execute traceroute` command from the subordinate unit CLI to display HA heartbeat IP addresses and the HA inter-VDOM link IP addresses. For example, use `execute ha manage 1` to connect to the subordinate unit CLI and then enter the following command to trace the route to an IP address on your network:

```
execute traceroute 172.20.20.10
traceroute to 172.20.20.10 (172.20.20.10), 32 hops max, 72 byte packets
 1 169.254.0.1 0 ms 0 ms 0 ms
 2 169.254.0.66 0 ms 0 ms 0 ms
 3 172.20.20.10 0 ms 0 ms 0 ms
```

Both HA heartbeat and data traffic are supported on the same FortiGate interface. All heartbeat communication takes place on a separate VDOM called `vsys_ha`. Heartbeat traffic uses a virtual interface called `port_ha` in the `vsys_ha` VDOM. Data and heartbeat traffic use the same physical interface, but they're logically separated into separate VDOMs.

Heartbeat packet Ethertypes

Normal IP packets are 802.3 packets that have an Ethernet type (Ethertype) field value of 0x0800. Ether type values other than 0x0800 are understood as level 2 frames rather than IP packets.

By default, HA heartbeat packets use the following Ethertypes:

- HA heartbeat packets for NAT/Route mode clusters use Ether type 0x8890. These packets are used by cluster units to find other cluster units and to verify the status of other cluster units while the cluster is operating. You can change the Ether type of these packets using the `ha-eth-type` option of the `config system ha` command.
- HA heartbeat packets for transparent mode clusters use Ether type 0x8891. These packets are used by cluster units to find other cluster units and to verify the status of other cluster units while the cluster is operating. You can change the Ether type of these packets using the `hc-eth-type` option of the `config system ha` command.
- HA telnet sessions between cluster units over HA heartbeat links use Ether type 0x8893. The telnet sessions allow an administrator to connect between FortiGates in the cluster using the `execute ha manage` command. You can change the Ether type of these packets using the `l2ep-eth-type` option of the `config system ha` command.

Because heartbeat packets are recognized as level 2 frames, the switches and routers on your heartbeat network that connect to heartbeat interfaces must be configured to allow them. If level2 frames are dropped by these network devices, heartbeat traffic will not be allowed between the cluster units.

Some third-party network equipment may use packets with these Ethertypes for other purposes. For example, Cisco N5K/Nexus switches use Ether type 0x8890 for some functions. When one of these switches receives Ether type 0x8890 packets from an attached cluster unit, the switch generates CRC errors and the packets are not forwarded. As a result, FortiGates connected with these switches cannot form a cluster.

In some cases, if the heartbeat interfaces are connected and configured so regular traffic flows but heartbeat traffic is not forwarded, you can change the configuration of the switch that connects the HA heartbeat interfaces to allow level2 frames with Ethertypes 0x8890, 0x8891, and 0x8893 to pass.

Alternatively, you can use the following CLI options to change the Ethertypes of the HA heartbeat packets:

```
config system ha
  set ha-eth-type <ha_ethertype_4-digit_hex>
  set hc-eth-type <hc_ethertype_4-digit_hex>
  set l2ep-eth-type <l2ep_ethertype_4-digit_hex>
end
```

For example, use the following command to change the Ether type of the HA heartbeat packets from 0x8890 to 0x8895 and to change the Ether type of HA Telnet session packets from 0x8891 to 0x889f:

```
config system ha
    set ha-eth-type 8895
    set l2ep-eth-type 889f
end
```

Modifying heartbeat timing

In an HA cluster, if a cluster unit CPU becomes very busy, the cluster unit may not be able to send heartbeat packets on time. If heartbeat packets are not sent on time other units in the cluster may think that the cluster unit has failed and the cluster will experience a failover.

A cluster unit CPU may become very busy if the cluster is subject to a syn flood attack, if network traffic is very heavy, or for other similar reasons. You can use the following CLI commands to configure how the cluster times HA heartbeat packets:

```
config system ha
    set hb-interval <interval_integer>
    set hb-lost-threshold <threshold_integer>
    set hello-holddown <holddown_integer>
end
```

Changing the lost heartbeat threshold

The lost heartbeat threshold is the number of consecutive heartbeat packets that are not received from another cluster unit before assuming that the cluster unit has failed. The default value is 6, meaning that if the 6 heartbeat packets are not received from a cluster unit then that cluster unit is considered to have failed. The range is 1 to 60 packets.

If the primary unit does not receive a heartbeat packet from a subordinate unit before the heartbeat threshold expires, the primary unit assumes that the subordinate unit has failed.

If a subordinate unit does not receive a heartbeat packet from the primary unit before the heartbeat threshold expires, the subordinate unit assumes that the primary unit has failed. The subordinate unit then begins negotiating to become the new primary unit.

The lower the `hb-lost-threshold` the faster a cluster responds when a unit fails. However, sometimes heartbeat packets may not be sent because a cluster unit is very busy. This can lead to a false positive failure detection. To reduce these false positives you can increase the `hb-lost-threshold`.

Use the following CLI command to increase the lost heartbeat threshold to 12:

```
config system ha
    set hb-lost-threshold 12
end
```

Changing the heartbeat interval

The heartbeat interval is the time between sending HA heartbeat packets. The heartbeat interval range is 1 to 20 (100*ms). The heartbeat interval default is 2 (200 ms).

A heartbeat interval of 2 means the time between heartbeat packets is 200 ms. Changing the heartbeat interval to 5 changes the time between heartbeat packets to 500 ms (5 * 100ms = 500ms).

The HA heartbeat packets consume more bandwidth if the heartbeat interval is short. But if the heartbeat interval is very long, the cluster is not as sensitive to topology and other network changes.

Use the following CLI command to increase the heartbeat interval to 10:

```
config system ha
    set hb-interval 10
end
```

The heartbeat interval combines with the lost heartbeat threshold to set how long a cluster unit waits before assuming that another cluster unit has failed and is no longer sending heartbeat packets. By default, if a cluster unit does not receive a heartbeat packet from a cluster unit for $6 * 200 = 1200$ milliseconds or 1.2 seconds the cluster unit assumes that the other cluster unit has failed.

You can increase both the heartbeat interval and the lost heartbeat threshold to reduce false positives. For example, increasing the heartbeat interval to 20 and the lost heartbeat threshold to 30 means a failure will be assumed if no heartbeat packets are received after $30 * 2000$ milliseconds = 60,000 milliseconds, or 60 seconds.

Use the following CLI command to increase the heartbeat interval to 20 and the lost heartbeat threshold to 30:

```
config system ha
    set hb-lost-threshold 30
    set hb-interval 20
end
```

Changing the time to wait in the hello state

The hello state hold-down time is the number of seconds that a cluster unit waits before changing from hello state to work state. After a failure or when starting up, cluster units operate in the hello state to send and receive heartbeat packets so that all the cluster units can find each other and form a cluster. A cluster unit should change from the hello state to work state after it finds all of the other FortiGate units to form a cluster with. If for some reason all cluster units cannot find each other during the hello state then some cluster units may be joining the cluster after it has formed. This can cause disruptions to the cluster and affect how it operates.

One reason for a delay in all of the cluster units joining the cluster could be the cluster units are located at different sites or if for some other reason communication is delayed between the heartbeat interfaces.

If cluster units are joining your cluster after it has started up or if it takes a while for units to join the cluster you can increase the time that the cluster units wait in the hello state. The hello state hold-down time range is 5 to 300 seconds. The hello state hold-down time default is 20 seconds.

Use the following CLI command to increase the time to wait in the hello state to 1 minute (60 seconds):

```
config system ha
    set hello-holddown 60
end
```

Enabling or disabling HA heartbeat encryption and authentication

You can enable HA heartbeat encryption and authentication to encrypt and authenticate HA heartbeat packets. HA heartbeat packets should be encrypted and authenticated if the cluster interfaces that send HA heartbeat packets are also connected to your networks.

If HA heartbeat packets are not encrypted the cluster password and changes to the cluster configuration could be exposed and an attacker may be able to sniff HA packets to get cluster information. Enabling HA heartbeat message authentication prevents an attacker from creating false HA heartbeat messages. False HA heartbeat messages could affect the stability of the cluster.

HA heartbeat encryption and authentication are disabled by default. Enabling HA encryption and authentication could reduce cluster performance. Use the following CLI command to enable HA heartbeat encryption and authentication.

```
config system ha
```



```
set authentication enable
set encryption enable
end
```

HA authentication and encryption uses AES-128 for encryption and SHA1 for authentication.

Heartbeat bandwidth requirements

The majority of the traffic processed by the HA heartbeat interface is session synchronization traffic. Other heartbeat interface traffic required to synchronize IPsec state/keys, routing tables, configuration changes, and so on is usually negligible.

The amount of traffic required for session synchronization depends on the connections per second (CPS) that the cluster is processing since only new sessions (and session table updates) need to be synchronized.

Another factor to consider is that if session pickup is enabled, traffic on the heartbeat interface surges during a failover or when a unit joins or re-joins the cluster. When one of these events happens, the whole session table needs to be synchronized. Lower bandwidth HA heartbeat interfaces may increase failover time if they can't handle the higher demand during these events.

You can also reduce the amount of heartbeat traffic by:

- Turning off session pickup if you don't need it,
- Configuring `session-pickup-delay` to reduce the number of sessions that are synchronized,
- Using the `session-sync-dev` option to move session synchronization traffic off of the heartbeat link.

See [Improving session synchronization performance on page 1](#) for details.

Cluster virtual MAC addresses

When a cluster is operating, the FGCP assigns virtual MAC addresses to each primary unit interface. HA uses virtual MAC addresses so that if a failover occurs, the new primary unit interfaces will have the same virtual MAC addresses and IP addresses as the failed primary unit. As a result, most network equipment would identify the new primary unit as the exact same device as the failed primary unit.

If the MAC addresses changed after a failover, the network would take longer to recover because all attached network devices would have to learn the new MAC addresses before they could communicate with the cluster.

If a cluster is operating in NAT/Route mode, the FGCP assigns a different virtual MAC address to each primary unit interface. VLAN subinterfaces are assigned the same virtual MAC address as the physical interface that the VLAN subinterface is added to. Redundant interfaces or 802.3ad aggregate interfaces are assigned the virtual MAC address of the first interface in the redundant or aggregate list.

If a cluster is operating in transparent mode, the FGCP assigns a virtual MAC address for the primary unit management IP address. Since you can connect to the management IP address from any interface, all of the FortiGate interfaces appear to have the same virtual MAC address.



A MAC address conflict can occur if two clusters are operating on the same network. See [Diagnosing packet loss with two FortiGate HA clusters in the same broadcast domain on page 1535](#) for more information.



Subordinate unit MAC addresses do not change. You can verify this by connecting to the subordinate unit CLI and using the `get hardware interface nic` command to display the MAC addresses of each FortiGate interface.



The MAC address of a reserved management interface is not changed to a virtual MAC address. Instead the reserved management interface keeps its original MAC address.

When the new primary unit is selected after a failover, the primary unit sends gratuitous ARP packets to update the devices connected to the cluster interfaces (usually layer-2 switches) with the virtual MAC address. Gratuitous ARP packets configure connected network devices to associate the cluster virtual MAC addresses and cluster IP address with primary unit physical interfaces and with the layer-2 switch physical interfaces. This is sometimes called using gratuitous ARP packets (sometimes called GARP packets) to train the network. The gratuitous ARP packets sent from the primary unit are intended to make sure that the layer-2 switch forwarding databases (FDBs) are updated as quickly as possible.

Sending gratuitous ARP packets is not required for routers and hosts on the network because the new primary unit will have the same MAC and IP addresses as the failed primary unit. However, since the new primary unit interfaces are connected to different switch interfaces than the failed primary unit, many network switches will update their FDBs more quickly after a failover if the new primary unit sends gratuitous ARP packets.

Changing how the primary unit sends gratuitous ARP packets after a failover

When a failover occurs it is important that the devices connected to the primary unit update their FDBs as quickly as possible to reestablish traffic forwarding.

Depending on your network configuration, you may be able to change the number of gratuitous ARP packets and the time interval between ARP packets to reduce the cluster failover time.

You cannot disable sending gratuitous ARP packets, but you can use the following command to change the number of packets that are sent. For example, enter the following command to send 20 gratuitous ARP packets:

```
config system ha
    set arps 20
end
```

You can use this command to configure the primary unit to send from 1 to 60 ARP packets. Usually you would not change the default setting of 5. In some cases, however, you might want to reduce the number of gratuitous ARP packets. For example, if your cluster has a large number of VLAN interfaces and virtual domains and because gratuitous ARP packets are broadcast, sending a higher number gratuitous ARP packets may generate a lot of network traffic. As long as the cluster still fails over successfully, you could reduce the number of gratuitous ARP packets that are sent to reduce the amount of traffic produced after a failover.

If failover is taking longer than expected, you may be able to reduce the failover time by increasing the number of gratuitous ARP packets sent.

You can also use the following command to change the time interval in seconds between gratuitous ARP packets. For example, enter the following command to change the time between ARP packets to 3 seconds:

```
config system ha
    set arps-interval 3
end
```

The time interval can be in the range of 1 to 20 seconds. The default is 8 seconds between gratuitous ARP packets. Normally you would not need to change the time interval. However, you could decrease the time to be able to send more packets in less time if your cluster takes a long time to failover.

There may also be a number of reasons to set the interval higher. For example, if your cluster has a large number of VLAN interfaces and virtual domains and because gratuitous ARP packets are broadcast, sending gratuitous

ARP packets may generate a lot of network traffic. As long as the cluster still fails over successfully you could increase the interval to reduce the amount of traffic produced after a failover.

For more information about gratuitous ARP packets see [RFC 826](#) and [RFC 3927](#).

Disabling gratuitous ARP packets after a failover

You can use the following command to turn off sending gratuitous ARP packets after a failover:

```
config system ha
    set gratuitous-arps disable
end
```

Sending gratuitous ARP packets is turned on by default.

In most cases you would want to send gratuitous ARP packets because its a reliable way for the cluster to notify the network to send traffic to the new primary unit. However, in some cases, sending gratuitous ARP packets may be less optimal. For example, if you have a cluster of FortiGates in transparent mode, after a failover the new primary unit will send gratuitous ARP packets to all of the addresses in its Forwarding Database (FDB). If the FDB has a large number of addresses it may take extra time to send all the packets and the sudden burst of traffic could disrupt the network.

If you choose to disable sending gratuitous ARP packets you must first enable the `link-failed-signal` setting. The cluster must have some way of informing attached network devices that a failover has occurred.

For more information about the `link-failed-signal` setting, see [Updating MAC forwarding tables when a link failover occurs on page 1553](#).

How the virtual MAC address is determined

The virtual MAC address is determined based on following formula:

$$00-09-0f-09- <group-id_hex> - (<vcluster_integer> + <idx>)$$

where

`<group-id_hex>` is the HA Group ID for the cluster converted to hexadecimal. The following table lists the virtual MAC address set for each group ID.

HA group ID in integer and hexadecimal format

Integer Group ID	Hexadecimal Group ID
0	00
1	01
2	02
3	03
4	04
...	...

Integer Group ID	Hexadecimal Group ID
10	0a
11	0b
...	...
63	3f
...	...
255	ff

`<vcluster_integer>` is 0 for virtual cluster 1 and 20 for virtual cluster 2. If virtual domains are not enabled, HA sets the virtual cluster to 1 and by default all interfaces are in the root virtual domain. Including virtual cluster and virtual domain factors in the virtual MAC address formula means that the same formula can be used whether or not virtual domains and virtual clustering is enabled.

`<idx>` is the index number of the interface. Interfaces are numbered from 0 to x (where x is the number of interfaces). Interfaces are numbered according to their has map order. See [Interface index and display order on page 1526](#). The first interface has an index of 0. The second interface in the list has an index of 1 and so on.



Only the `<idx>` part of the virtual MAC address is different for each interface. The `<vcluster_integer>` would be different for different interfaces if multiple VDOMs have been added.



Between FortiOS releases interface indexing may change so the virtual MAC addresses assigned to individual FortiGate interfaces may also change.

Example virtual MAC addresses

An HA cluster with HA group ID unchanged (default=0) and virtual domains not enabled would have the following virtual MAC addresses for interfaces port1 to port12:

- port1 virtual MAC: 00-09-0f-09-00-00
- port10 virtual MAC: 00-09-0f-09-00-01
- port2 virtual MAC: 00-09-0f-09-00-02
- port3 virtual MAC: 00-09-0f-09-00-03
- port4 virtual MAC: 00-09-0f-09-00-04
- port5 virtual MAC: 00-09-0f-09-00-05
- port6 virtual MAC: 00-09-0f-09-00-06
- port7 virtual MAC: 00-09-0f-09-00-07
- port8 virtual MAC: 00-09-0f-09-00-08
- port9 virtual MAC: 00-09-0f-09-00-

- port11 virtual MAC: 00-09-0f-09-00-0a
- port12 virtual MAC: 00-09-0f-09-00-0b

If the group ID is changed to 34 these virtual MAC addresses change to:

- port1 virtual MAC: 00-09-0f-09-22-00
- port3 virtual MAC: 00-09-0f-09-22-03
- port4 virtual MAC: 00-09-0f-09-22-04
- port5 virtual MAC: 00-09-0f-09-22-05
- port6 virtual MAC: 00-09-0f-09-22-06
- port7 virtual MAC: 00-09-0f-09-22-07
- port8 virtual MAC: 00-09-0f-09-22-08
- port9 virtual MAC: 00-09-0f-09-22-
- port11 virtual MAC: 00-09-0f-09-22-0a
- port12 virtual MAC: 00-09-0f-09-22-0b
- port10 virtual MAC: 00-09-0f-09-22-01
- port2 virtual MAC: 00-09-0f-09-22-02

A cluster with virtual domains enabled where the HA group ID has been changed to 23, port5 and port 6 are in the root virtual domain (which is in virtual cluster1), and port7 and port8 are in the vdom_1 virtual domain (which is in virtual cluster 2) would have the following virtual MAC addresses:

- port5 interface virtual MAC: 00-09-0f-09-23-05
- port6 interface virtual MAC: 00-09-0f-09-23-06
- port7 interface virtual MAC: 00-09-0f-09-23-27
- port8 interface virtual MAC: 00-09-0f-09-23-28

Displaying the virtual MAC address

Every FortiGate physical interface has two MAC addresses: the current hardware address and the permanent hardware address. The permanent hardware address cannot be changed, it is the actual MAC address of the interface hardware. The current hardware address can be changed. The current hardware address is the address seen by the network. For a FortiGate not operating in HA, you can use the following command to change the current hardware address of the port1 interface:

```
config system interface
  edit port1
    set macaddr <mac_address>
  end
end
```

For an operating cluster, the current hardware address of each cluster unit interface is changed to the HA virtual MAC address by the FGCP. The `macaddr` option is not available for a functioning cluster. You cannot change an interface MAC address and you cannot view MAC addresses from the `system interface` CLI command.

You can use the `get hardware nic <interface_name_str>` command to display both MAC addresses for any FortiGate interface. This command displays hardware information for the specified interface. Depending on their hardware configuration, this command may display different information for different interfaces. You can use this command to display the current hardware address as `Current_HWaddr` and the permanent hardware address as `Permanent_HWaddr`. For some interfaces the current hardware address is displayed as `MAC`. The command displays a great deal of information about the interface so you may have to scroll the output to find the hardware addresses.



You can also use the `diagnose hardware deviceinfo nic <interface_str>` command to display both MAC addresses for any FortiGate interface.

Before HA configuration the current and permanent hardware addresses are the same. For example for one of the units in Cluster_1:

```
FGT60B3907503171 # get hardware nic internal
.
.
.
MAC: 02:09:0f:78:18:c9
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

During HA operation the current hardware address becomes the HA virtual MAC address, for example for the units in Cluster_1:

```
FGT60B3907503171 # get hardware nic internal
.
.
.
MAC: 00:09:0f:09:00:02
Permanent_HWaddr: 02:09:0f:78:18:c9
.
.
.
```

The following command output for Cluster_2 shows the same current hardware address for port1 as for the internal interface of Cluster_2, indicating a MAC address conflict.

```
FG300A2904500238 # get hardware nic port1
.
.
.
MAC: 00:09:0f:09:00:02
Permanent_HWaddr: 00:09:0F:85:40:FD
.
.
.
```

Diagnosing packet loss with two FortiGate HA clusters in the same broadcast domain

A network may experience packet loss when two FortiGate HA clusters have been deployed in the same broadcast domain. Deploying two HA clusters in the same broadcast domain can result in packet loss because of MAC address conflicts. The packet loss can be diagnosed by pinging from one cluster to the other or by pinging both of the clusters from a device within the broadcast domain. You can resolve the MAC address conflict by changing the HA Group ID configuration of the two clusters. The HA Group ID is sometimes also called the Cluster ID.

This section describes a topology that can result in packet loss, how to determine if packets are being lost, and how to correct the problem by changing the HA Group ID.



Packet loss on a network can also be caused by IP address conflicts. Finding and fixing IP address conflicts can be difficult. However, if you are experiencing packet loss and your network contains two FortiGate HA clusters you can use the information in this article to eliminate one possible source of packet loss.

Changing the HA group ID to avoid MAC address conflicts

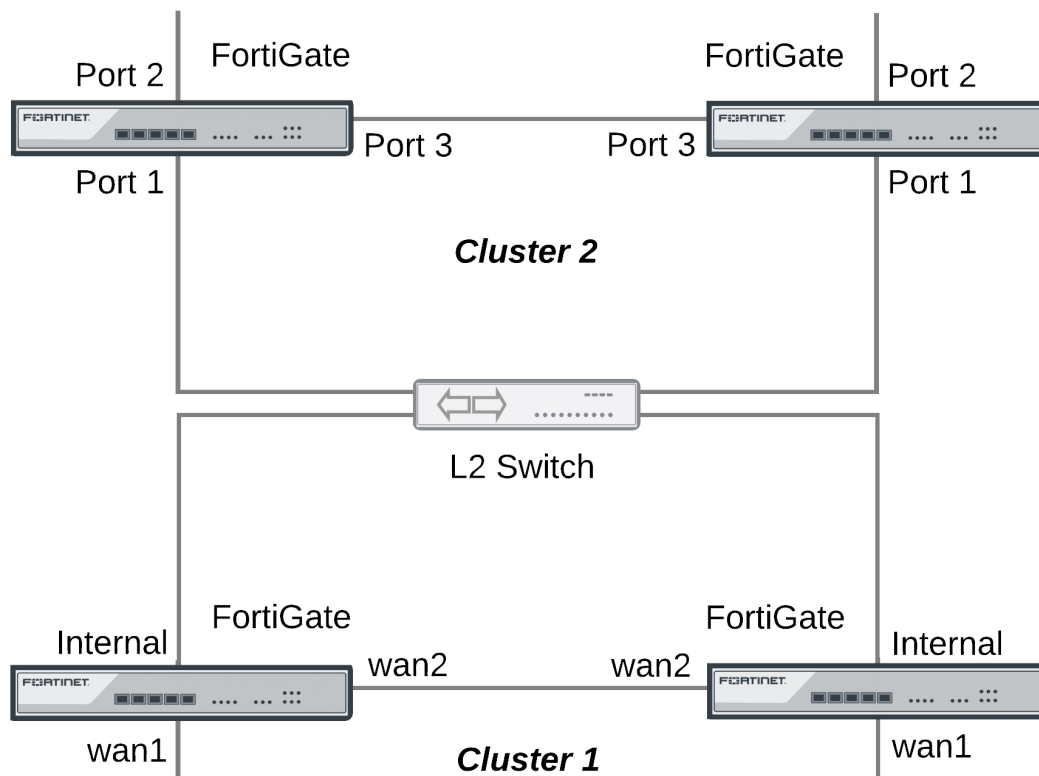
Change the Group ID to change the virtual MAC address of all cluster interfaces. You can change the Group ID from the FortiGate CLI using the following command:

```
config system ha
    set group-id <id_integer>
end
```

Example topology

The topology below shows two clusters. The Cluster_1 internal interfaces and the Cluster_2 port 1 interfaces are both connected to the same broadcast domain. In this topology the broadcast domain could be an internal network. Both clusters could also be connected to the internet or to different networks.

Example HA topology with possible MAC address conflicts



Ping testing for packet loss

If the network is experiencing packet loss, it is possible that you will not notice a problem unless you are constantly pinging both HA clusters. During normal operation of the network you also might not notice packet loss because the loss rate may not be severe enough to timeout TCP sessions. Also many common types of TCP traffic, such as web browsing, may not be greatly affected by packet loss. However, packet loss can have a significant effect on real time protocols that deliver audio and video data.

To test for packet loss you can set up two constant ping sessions, one to each cluster. If packet loss is occurring the two ping sessions should show alternating replies and timeouts from each cluster.

Cluster_1	Cluster_2
reply	timeout
reply	timeout
reply	timeout
timeout	reply
timeout	reply
reply	timeout
reply	timeout
timeout	reply
timeout	reply
timeout	reply
timeout	reply

Viewing MAC address conflicts on attached switches

If two HA clusters with the same virtual MAC address are connected to the same broadcast domain (L2 switch or hub), the MAC address will conflict and bounce between the two clusters. This example Cisco switch MAC address table shows the MAC address flapping between different interfaces (1/0/1 and 1/0/4).

```
1 0009.0f09.0002 DYNAMIC Gi1/0/1
1 0009.0f09.0002 DYNAMIC Gi1/0/4
```

Synchronizing the configuration

The FGCP uses a combination of incremental and periodic synchronization to make sure that the configuration of all cluster units is synchronized to that of the primary unit.

The following settings are not synchronized between cluster units:

- HA override.
- HA device priority.

- The virtual cluster priority.
- The FortiGate host name.
- The HA priority setting for a ping server (or dead gateway detection) configuration.
- The system interface settings of the HA reserved management interface.
- The HA default route for the reserved management interface, set using the `ha-mgmt-interface-gateway` option of the `config system ha` command.

The primary unit synchronizes all other configuration settings, including the other HA configuration settings.

All synchronization activity takes place over the HA heartbeat link using TCP/703 and UDP/703 packets.

Disabling automatic configuration synchronization

In some cases you may want to use the following command to disable automatic synchronization of the primary unit configuration to all cluster units.

```
config system ha
    set sync-config disable
end
```

When this option is disabled the cluster no longer synchronizes configuration changes. If a device failure occurs, the new primary unit may not have the same configuration as the failed primary unit. As a result, the new primary unit may process sessions differently or may not function on the network in the same way.

In most cases you should not disable automatic configuration synchronization. However, if you have disabled this feature you can use the `execute ha synchronize` command to manually synchronize a subordinate unit's configuration to that of the primary unit.

You must enter `execute ha synchronize` commands from the subordinate unit that you want to synchronize with the primary unit. Use the `execute ha manage` command to access a subordinate unit CLI.

For example, to access the first subordinate unit and force a synchronization at any time, even if automatic synchronization is disabled enter:

```
execute ha manage 0
execute ha synchronize start
```

You can use the following command to stop a synchronization that is in progress.

```
execute ha synchronize stop
```

Incremental synchronization

When you log into the cluster GUI or CLI to make configuration changes, you are actually logging into the primary unit. All of your configuration changes are first made to the primary unit. Incremental synchronization then immediately synchronizes these changes to all of the subordinate units.

When you log into a subordinate unit CLI (for example using `execute ha manage`) all of the configuration changes that you make to the subordinate unit are also immediately synchronized to all cluster units, including the primary unit, using the same process.

Incremental synchronization also synchronizes other dynamic configuration information such as the DHCP server address lease database, routing table updates, IPsec SAs, MAC address tables, and so on. See [FortiGate HA compatibility with DHCP and PPPoE on page 1397](#) for more information about DHCP server address lease synchronization and [Synchronizing kernel routing tables on page 1546](#) for information about routing table updates.

Whenever a change is made to a cluster unit configuration, incremental synchronization sends the same configuration change to all other cluster units over the HA heartbeat link. An HA synchronization process running on the each cluster unit receives the configuration change and applies it to the cluster unit. The HA synchronization process makes the configuration change by entering a CLI command that appears to be entered by the administrator who made the configuration change in the first place.

Synchronization takes place silently, and no log messages are recorded about the synchronization activity. However, log messages can be recorded by the cluster units when the synchronization process enters CLI commands. You can see these log messages on the subordinate units if you enable event logging and set the minimum severity level to **Information** and then check the event log messages written by the cluster units when you make a configuration change.

You can also see these log messages on the primary unit if you make configuration changes from a subordinate unit.

Periodic synchronization

Incremental synchronization makes sure that as an administrator makes configuration changes, the configurations of all cluster units remain the same. However, a number of factors could cause one or more cluster units to go out of sync with the primary unit. For example, if you add a new unit to a functioning cluster, the configuration of this new unit will not match the configuration of the other cluster units. Its not practical to use incremental synchronization to change the configuration of the new unit.

Periodic synchronization is a mechanism that looks for synchronization problems and fixes them. Every minute the cluster compares the configuration file checksum of the primary unit with the configuration file checksums of each of the subordinate units. If all subordinate unit checksums are the same as the primary unit checksum, all cluster units are considered synchronized.

If one or more of the subordinate unit checksums is not the same as the primary unit checksum, the subordinate unit configuration is considered out of sync with the primary unit. The checksum of the out of sync subordinate unit is checked again every 15 seconds. This re-checking occurs in case the configurations are out of sync because an incremental configuration sequence has not completed. If the checksums do not match after 5 checks the subordinate unit that is out of sync retrieves the configuration from the primary unit. The subordinate unit then reloads its configuration and resumes operating as a subordinate unit with the same configuration as the primary unit.

The configuration of the subordinate unit is reset in this way because when a subordinate unit configuration gets out of sync with the primary unit configuration there is no efficient way to determine what the configuration differences are and to correct them. Resetting the subordinate unit configuration becomes the most efficient way to resynchronize the subordinate unit.

Synchronization requires that all cluster units run the same FortiOS firmware build. If some cluster units are running different firmware builds, then unstable cluster operation may occur and the cluster units may not be able to synchronize correctly.



Re-installing the firmware build running on the primary unit forces the primary unit to upgrade all cluster units to the same firmware build.

Console messages when configuration synchronization succeeds

When a cluster first forms, or when a new unit is added to a cluster as a subordinate unit, the following messages appear on the CLI console to indicate that the unit joined the cluster and had its configuring synchronized with the

primary unit.

```
slave's configuration is not in sync with master's, sequence:0
slave's configuration is not in sync with master's, sequence:1
  slave's configuration is not in sync with master's, sequence:2
slave's configuration is not in sync with master's, sequence:3
  slave's configuration is not in sync with master's, sequence:4
  slave starts to sync with master
logout all admin users
slave succeeded to sync with master
```

Console messages when configuration synchronization fails

If you connect to the console of a subordinate unit that is out of synchronization with the primary unit, messages similar to the following are displayed.

```
slave is not in sync with master, sequence:0. (type 0x3)
slave is not in sync with master, sequence:1. (type 0x3)
slave is not in sync with master, sequence:2. (type 0x3)
slave is not in sync with master, sequence:3. (type 0x3)
slave is not in sync with master, sequence:4. (type 0x3)
global compared not matched
```

If synchronization problems occur the console message sequence may be repeated over and over again. The messages all include a type value (in the example `type 0x3`). The type value can help Fortinet Support diagnose the synchronization problem.

HA out of sync object messages and the configuration objects that they reference

Out of Sync Message	Configuration Object
HA_SYNC_SETTING_CONFIGURATION = 0x03	/data/config
HA_SYNC_SETTING_AV = 0x10	
HA_SYNC_SETTING_VIR_DB = 0x11	/etc/vir
HA_SYNC_SETTING_SHARED_LIB = 0x12	/data/lib/libav.so
HA_SYNC_SETTING_SCAN_UNIT = 0x13	/bin/scanunitd
HA_SYNC_SETTING_IMAP_PRXY = 0x14	/bin/imapd
HA_SYNC_SETTING_SMTP_PRXY = 0x15	/bin/smtp
HA_SYNC_SETTING_POP3_PRXY = 0x16	/bin/pop3
HA_SYNC_SETTING_HTTP_PRXY = 0x17	/bin/thttp
HA_SYNC_SETTING_FTP_PRXY = 0x18	/bin/ftpd
HA_SYNC_SETTING_FCNI = 0x19	/etc/fcni.dat

Out of Sync Message	Configuration Object
HA_SYNC_SETTING_FDNI = 0x1a	/etc/fdnserver.dat
HA_SYNC_SETTING_FSCI = 0x1b	/etc/sci.dat
HA_SYNC_SETTING_FSAE = 0x1c	/etc/fsae_adgrp.cache
HA_SYNC_SETTING_IDS = 0x20	/etc/ids.rules
HA_SYNC_SETTING_IDSUSER_RULES = 0x21	/etc/idsuser.rules
HA_SYNC_SETTING_IDSCUSTOM = 0x22	
HA_SYNC_SETTING_IDS_MONITOR = 0x23	/bin/ipsmonitor
HA_SYNC_SETTING_IDS_SENSOR = 0x24	/bin/ipsengine
HA_SYNC_SETTING_NIDS_LIB = 0x25	/data/lib/libips.so
HA_SYNC_SETTING_WEBLISTS = 0x30	
HA_SYNC_SETTING_CONTENTFILTER = 0x31	/data/cmdb/webfilter.bword
HA_SYNC_SETTING_URLFILTER = 0x32	/data/cmdb/webfilter.urlfilter
HA_SYNC_SETTING_FTGD_OVRD = 0x33	/data/cmdb/webfilter.fgtd-ovrd
HA_SYNC_SETTING_FTGD_LRATING = 0x34	/data/cmdb/webfilter.fgtd-ovrd
HA_SYNC_SETTING_EMAILLISTS = 0x40	
HA_SYNC_SETTING_EMAILCONTENT = 0x41	/data/cmdb/spamfilter.bword
HA_SYNC_SETTING_EMAILBWLIST = 0x42	/data/cmdb/spamfilter.emailbwl
HA_SYNC_SETTING_IPBWL = 0x43	/data/cmdb/spamfilter.ipbwl
HA_SYNC_SETTING_MHEADER = 0x44	/data/cmdb/spamfilter.mheader
HA_SYNC_SETTING_RBL = 0x45	/data/cmdb/spamfilter.rbl
HA_SYNC_SETTING_CERT_CONF = 0x50	/etc/cert/cert.conf
HA_SYNC_SETTING_CERT_CA = 0x51	/etc/cert/ca
HA_SYNC_SETTING_CERT_LOCAL = 0x52	/etc/cert/local
HA_SYNC_SETTING_CERT_CRL = 0x53	/etc/cert/crl

Out of Sync Message	Configuration Object
HA_SYNC_SETTING_DB_VER = 0x55	
HA_GET_DETAIL_CSUM = 0x71	
HA_SYNC_CC_SIG = 0x75	/etc/cc_sig.dat
HA_SYNC_CC_OP = 0x76	/etc/cc_op
HA_SYNC_CC_MAIN = 0x77	/etc/cc_main
HA_SYNC_FTGD_CAT_LIST = 0x7a	/migadmin/webfilter/ublock/ftgd/data/

Comparing checksums of cluster units

You can use the `diagnose sys ha checksum show` command to compare the configuration checksums of all cluster units. The output of this command shows checksums labeled `global` and `all` as well as checksums for each of the VDOMs including the `root` VDOM. The `get system ha-nonsync-csum` command can be used to display similar information; however, this command is intended to be used by FortiManager.

The primary unit and subordinate unit checksums should be the same. If they are not you can use the `execute ha synchronize start` command to force a synchronization.

The following command output is for the primary unit of a cluster that does not have multiple VDOMs enabled:

```
diagnose sys ha checksum show
is_manage_master()=1, is_root_master()=1
debugzone
global: a0 7f a7 ff ac 00 d5 b6 82 37 cc 13 3e 0b 9b 77
root: 43 72 47 68 7b da 81 17 c8 f5 10 dd fd 6b e9 57
all: c5 90 ed 22 24 3e 96 06 44 35 b6 63 7c 84 88 d5

checksum
global: a0 7f a7 ff ac 00 d5 b6 82 37 cc 13 3e 0b 9b 77
root: 43 72 47 68 7b da 81 17 c8 f5 10 dd fd 6b e9 57
all: c5 90 ed 22 24 3e 96 06 44 35 b6 63 7c 84 88 d5
```

The following command output is for a subordinate unit of the same cluster:

```
diagnose sys ha checksum show
is_manage_master()=0, is_root_master()=0
debugzone
global: a0 7f a7 ff ac 00 d5 b6 82 37 cc 13 3e 0b 9b 77
root: 43 72 47 68 7b da 81 17 c8 f5 10 dd fd 6b e9 57
all: c5 90 ed 22 24 3e 96 06 44 35 b6 63 7c 84 88 d5

checksum
global: a0 7f a7 ff ac 00 d5 b6 82 37 cc 13 3e 0b 9b 77
root: 43 72 47 68 7b da 81 17 c8 f5 10 dd fd 6b e9 57
all: c5 90 ed 22 24 3e 96 06 44 35 b6 63 7c 84 88 d5
```

The following example shows using this command for the primary unit of a cluster with multiple VDOMs. Two VDOMs have been added named `test` and `Eng_vdm`.

From the primary unit:

```

config global
  diagnose sys ha checksum show
  is_manage_master()=1, is_root_master()=1
  debugzone
  global: 65 75 88 97 2d 58 1b bf 38 d3 3d 52 5b 0e 30 a9
  test: a5 16 34 8c 7a 46 d6 a4 1e 1f c8 64 ec 1b 53 fe
  root: 3c 12 45 98 69 f2 d8 08 24 cf 02 ea 71 57 a7 01
  Eng_vdm: 64 51 7c 58 97 79 b1 b3 b3 ed 5c ec cd 07 74 09
  all: 30 68 77 82 a1 5d 13 99 d1 42 a3 2f 9f b9 15 53

  checksum
  global: 65 75 88 97 2d 58 1b bf 38 d3 3d 52 5b 0e 30 a9
  test: a5 16 34 8c 7a 46 d6 a4 1e 1f c8 64 ec 1b 53 fe
  root: 3c 12 45 98 69 f2 d8 08 24 cf 02 ea 71 57 a7 01
  Eng_vdm: 64 51 7c 58 97 79 b1 b3 b3 ed 5c ec cd 07 74 09
  all: 30 68 77 82 a1 5d 13 99 d1 42 a3 2f 9f b9 15 53

```

From the subordinate unit:

```

config global
  diagnose sys ha checksum show
  is_manage_master()=0, is_root_master()=0
  debugzone
  global: 65 75 88 97 2d 58 1b bf 38 d3 3d 52 5b 0e 30 a9
  test: a5 16 34 8c 7a 46 d6 a4 1e 1f c8 64 ec 1b 53 fe
  root: 3c 12 45 98 69 f2 d8 08 24 cf 02 ea 71 57 a7 01
  Eng_vdm: 64 51 7c 58 97 79 b1 b3 b3 ed 5c ec cd 07 74 09
  all: 30 68 77 82 a1 5d 13 99 d1 42 a3 2f 9f b9 15 53

  checksum
  global: 65 75 88 97 2d 58 1b bf 38 d3 3d 52 5b 0e 30 a9
  test: a5 16 34 8c 7a 46 d6 a4 1e 1f c8 64 ec 1b 53 fe
  root: 3c 12 45 98 69 f2 d8 08 24 cf 02 ea 71 57 a7 01
  Eng_vdm: 64 51 7c 58 97 79 b1 b3 b3 ed 5c ec cd 07 74 09
  all: 30 68 77 82 a1 5d 13 99 d1 42 a3 2f 9f b9 15 53

```

How to diagnose HA out of sync messages

This section describes how to use the `diagnose sys ha checksum show` and `diagnose debug` commands to diagnose the cause of HA out of sync messages.

If HA synchronization is not successful, use the following procedures on each cluster unit to find the cause.

To determine why HA synchronization does not occur

1. Connect to each cluster unit CLI by connected to the console port.
2. Enter the following commands to enable debugging and display HA out of sync messages.

```

diagnose debug enable
diagnose debug console timestamp enable
diagnose debug application hataalk -1
diagnose debug application hasync -1

```

Collect the console output and compare the out of sync messages with the information in the table [HA out of sync object messages and the configuration objects that they reference on page 1540](#).

3. Enter the following commands to turn off debugging.

```

diagnose debug disable
diagnose debug reset

```

To determine what part of the configuration is causing the problem

If the previous procedure displays messages that include sync object 0x30 (for example, `HA_SYNC_SETTING_CONFIGURATION = 0x03`) there is a synchronization problem with the configuration. Use the following steps to determine the part of the configuration that is causing the problem.

If your cluster consists of two cluster units, use this procedure to capture the configuration checksums for each unit. If your cluster consists of more than two cluster units, repeat this procedure for all cluster units that returned messages that include 0x30 sync object messages.

1. Connect to each cluster unit CLI by connected to the console port.
2. Enter the following command to turn on terminal capture
`diagnose debug enable`
3. Enter the following command to stop HA synchronization.
`execute ha sync stop`
4. Enter the following command to display configuration checksums.
`diagnose sys ha checksum show global`
5. Copy the output to a text file.
6. Repeat for all affected units.
7. Compare the text file from the primary unit with the text file from each cluster unit to find the checksums that do not match.

You can use a diff function to compare text files.

8. Repeat for the root VDOM:
`diagnose sys ha checksum show root`
9. Repeat for all VDOMS (if multiple VDOM configuration is enabled):
`diagnose sys ha checksum show <vdom-name>`
10. You can also use the `grep` option to just display checksums for parts of the configuration.
For example to display system related configuration checksums in the root VDOM or log-related checksums in the global configuration:

```
diagnose sys ha checksum root | grep system
diagnose sys ha checksum global | grep log
```

Generally it is the first non-matching checksum that is the cause of the synchronization problem.

11. Attempt to remove/change the part of the configuration that is causing the problem. You can do this by making configuration changes from the primary unit or subordinate unit CLI.
12. Enter the following commands to start HA configuration and stop debugging:

```
execute ha sync start
diagnose debug disable
diagnose debug reset
```

Recalculating the checksums to resolve out of sync messages

Sometimes an error can occur when checksums are being calculated by the cluster. As a result of this calculation error the CLI console could display out of sync error messages even though the cluster is otherwise operating normally. You can also sometimes see checksum calculation errors in `diagnose sys ha checksum`

command output when the checksums listed in the `debugzone` output don't match the checksums in the `checksum` part of the output.

One solution to this problem could be to re-calculate the checksums. The re-calculated checksums should match and the out of sync error messages should stop appearing.

You can use the following command to re-calculate HA checksums:

```
diagnose sys ha checksum recalculate [<vdom-name> | global]
```

Just entering the command without options recalculates all checksums. You can specify a VDOM name to just recalculate the checksums for that VDOM. You can also enter `global` to recalculate the global checksum.

Determining what is causing a configuration synchronization problem

There are twenty-five FortiOS modules that have their configurations synchronized. It can be difficult to find the cause of a synchronization problem with so much data to analyze. You can use the following diagnose commands to more easily find modules that may be causing synchronization problems.

```
diagnose sys ha hasync-stats {all | most-recent [<seconds>] | by object [<number>]}
```

`all` displays the synchronization activity for all modules that happened since the hasync process started running (usually this would be since the cluster started-up).

`most-recent [<seconds>]` displays the most recently occurring synchronization events. You can include a time in seconds to display recent events that occurred during the time interval. If you don't include the number of seconds, the command displays the most recent events in the last 5 seconds. This option can be used to determine the module or modules that are currently synchronizing or attempting to synchronize. If no modules are currently synchronizing, the command just displays the most recent synchronization events.

`by-object [<number>]` displays the synchronization activity of a specific module, where `<number>` is the module number in the range 1 to 25. To display a list of all 25 modules and their numbers enter:

```
diagnose sys ha hasync-stats by-object ?
```

To display the most recent activity, enter:

```
diagnose sys ha hasync-stats most-recent 10
```

```
current-time/jiffies=2018-03-28 13:01:42/1148242:
```

```
hasync-obj=2(arp):
```

```
  epoll_handler=1(ev_arp_handler): start=1522256500.354400(2018-03-28 13:01:40),
end=1522256500.354406(2018-03-28 13:01:40), total=0.000006/1699
```

```
hasync-obj=5(config):
```

```
  timer=0(check_sync_status), add=1141764(2018-03-28 13:01:26), expire=1142764
(2018-03-28 13:01:36), end=1142764(2018-03-28 13:01:36), del=0(), total_call=1143
```

```
hasync-obj=8(time):
```

```
  obj_handler=0(packet): start=1522256497.851550(2018-03-28 13:01:37), end-
d=1522256497.851570(2018-03-28 13:01:37), total=0.000020/381
```

```
  timer=0(sync_timer), add=1140106(2018-03-28 13:01:10), expire=1143106(2018-03-
28 13:01:40), end=1143106(2018-03-28 13:01:40), del=0(), total_call=381
```

```
hasync-obj=21(hastats):
```

```
  obj_handler=0(packet): start=1522256499.760934(2018-03-28 13:01:39), end-
d=1522256499.760936(2018-03-28 13:01:39), total=0.000002/2285
```

```
  timer=0(hastats_timer), add=1142556(2018-03-28 13:01:34), expire=1143056(2018-
03-28 13:01:39), end=1143056(2018-03-28 13:01:39), del=0(), total_call=2286
```


The last few lines of this output shows activity with the `hastats` module, which is module 21. You can use the following command to see more information about synchronization activity with this module:

```
diagnose sys ha hasync-stats by-object 21
```

Synchronizing kernel routing tables

In a functioning cluster, the primary unit keeps all subordinate unit kernel routing tables (also called the forwarding information base FIB) up to date and synchronized with the primary unit. All synchronization activity takes place over the HA heartbeat link using TCP/703 and UDP/703 packets. After a failover, because of these routing table updates the new primary unit does not have to populate its kernel routing table before being able to route traffic. This gives the new primary unit time to rebuild its regular routing table after a failover.

Use the following command to view the regular routing table. This table contains all of the configured routes and routes acquired from dynamic routing protocols and so on. This routing table is not synchronized. On subordinate units this command will not produce the same output as on the primary unit.

```
get router info routing-table
```

Use the following command to view the kernel routing table (FIB). This is the list of resolved routes actually being used by the FortiOS kernel. The output of this command should be the same on the primary unit and the subordinate units.

```
get router info kernel
```

This section describes how clusters handle dynamic routing failover and also describes how to use CLI commands to control the timing of routing table updates of the subordinate unit routing tables from the primary unit.

Controlling how the FGCP synchronizes kernel routing table updates

You can use the following commands to control some of the timing settings that the FGCP uses when synchronizing routing updates from the primary unit to subordinate units and maintaining routes on the primary unit after a failover.

```
config system ha
  set route-hold <hold_integer>
  set route-ttl <ttl_integer>
  set route-wait <wait_integer>
end
```

Change how long routes stay in a cluster unit routing table

Change the `route-ttl` time to control how long routes remain in a cluster unit routing table. The time to live range is 5 to 3600 seconds. The default time to live is 10 seconds.

The time to live controls how long routes remain active in a cluster unit routing table after the cluster unit becomes a primary unit. To maintain communication sessions after a cluster unit becomes a primary unit, routes remain active in the routing table for the route time to live while the new primary unit acquires new routes.

By default, `route-ttl` is set to 10 which may mean that only a few routes will remain in the routing table after a failover. Normally keeping `route-ttl` to 10 or reducing the value to 5 is acceptable because acquiring new routes usually occurs very quickly, especially if graceful restart is enabled, so only a minor delay is caused by acquiring new routes.

If the primary unit needs to acquire a very large number of routes, or if for other reasons, there is a delay in acquiring all routes, the primary unit may not be able to maintain all communication sessions.

You can increase the route time to live if you find that communication sessions are lost after a failover so that the primary unit can use synchronized routes that are already in the routing table, instead of waiting to acquire new routes.

Change the time between routing updates

Change the `route-hold` time to change the time that the primary unit waits between sending routing table updates to subordinate units. The route hold range is 0 to 3600 seconds. The default route hold time is 10 seconds.

To avoid flooding routing table updates to subordinate units, set `route-hold` to a relatively long time to prevent subsequent updates from occurring too quickly. Flooding routing table updates can affect cluster performance if a great deal of routing information is synchronized between cluster units. Increasing the time between updates means that this data exchange will not have to happen so often.

The `route-hold` time should be coordinated with the `route-wait` time.

Change the time the primary unit waits after receiving a routing update

Change the `route-wait` time to change how long the primary unit waits after receiving routing updates before sending the updates to the subordinate units. For quick routing table updates to occur, set `route-wait` to a relatively short time so that the primary unit does not hold routing table changes for too long before updating the subordinate units.

The `route-wait` range is 0 to 3600 seconds. The default `route-wait` is 0 seconds.

Normally, because the `route-wait` time is 0 seconds the primary unit sends routing table updates to the subordinate units every time its routing table changes.

Once a routing table update is sent, the primary unit waits the `route-hold` time before sending the next update.

Usually routing table updates are periodic and sporadic. Subordinate units should receive these changes as soon as possible so `route-wait` is set to 0 seconds. `route-hold` can be set to a relatively long time because normally the next route update would not occur for a while.

In some cases, routing table updates can occur in bursts. A large burst of routing table updates can occur if a router or a link on a network fails or changes. When a burst of routing table updates occurs, there is a potential that the primary unit could flood the subordinate units with routing table updates. Flooding routing table updates can affect cluster performance if a great deal of routing information is synchronized between cluster units. Setting `route-wait` to a longer time reduces the frequency of additional updates and prevents flooding of routing table updates from occurring.

Configuring graceful restart for dynamic routing failover

When an HA failover occurs, neighbor routers will detect that the cluster has failed and remove it from the network until the routing topology stabilizes. During that time the routers may stop sending IP packets to the cluster and communication sessions that would normally be processed by the cluster may time out or be dropped. Also the new primary unit will not receive routing updates and so will not be able to build and maintain its routing database.

You can solve this problem by configuring graceful restart for the dynamic routing protocols that you are using. This section describes configuring graceful restart for OSPF and BGP.

To support graceful restart you should make sure the new primary unit keeps its synchronized routing data long enough to acquire new routing data. You should also increase the HA route time to live, route wait, and route hold values to 60 using the following CLI command:

```
config system ha
    set route-ttl 60
    set route-wait 60
    set route-hold 60
end
```

Graceful OSPF restart

You can configure graceful restart (also called nonstop forwarding (NSF) as described in [RFC3623](#) (Graceful OSPF Restart) to solve the problem of dynamic routing failover. If graceful restart is enabled on neighbor routers, they will keep sending packets to the cluster following the HA failover instead of removing it from the network. The neighboring routers assume that the cluster is experiencing a graceful restart.

After the failover, the new primary unit can continue to process communication sessions using the synchronized routing data received from the failed primary unit before the failover. This gives the new primary unit time to update its routing table after the failover.

You can use the following commands to enable graceful restart or NSF on Cisco routers:

```
router ospf 1
    log-adjacency-changes
    nsf ietf helper strict-lsa-checking
```

If the cluster is running OSPF, use the following command to enable graceful restart for OSPF:

```
config router ospf
    set restart-mode graceful-restart
end
```

Graceful BGP restart

If the cluster is running BGP only the primary unit keeps BGP peering connections. When a failover occurs, the BGP peering needs to be reestablished. This will happen if you enable BGP graceful restart which causes the adjacent routers to keep the routes active while the BGP peering is restarted by the new primary unit.



Enabling BGP graceful restart causes the FortiGate BGP process to restart which can temporarily disrupt traffic through the cluster. So normally you should wait for a quiet time or a maintenance period to enable BGP graceful restart.

Use the following command to enable graceful restart for BGP and set some graceful restart options.

```
config router bgp
    set graceful-restart enable
    set graceful-restart-time 120
    set graceful-stalepath-time 360
    set graceful-update-delay 120
end
```

Notifying BGP neighbors when graceful restart is enabled

You can add BGP neighbors and configure the cluster unit to notify these neighbors that it supports graceful restart.

```
config router bgp
```

```
config neighbor
  edit <neighbor_address_Ipv4>
    set capability-graceful-restart enable
  end
end
```

Bidirectional Forwarding Detection (BFD) enabled BGP graceful restart

You can add a BFD enabled BGP neighbor as a static BFD neighbor using the following command. This example shows how to add a BFD neighbor with IP address 172.20.121.23 that is on the network connected to port4:

```
config router bfd
  config neighbor
    edit 172.20.121.23
      set port4
    end
  end
end
```

The FGCP supports graceful restart of BFD enabled BGP neighbors. The `config router bfd` command is needed as the BGP auto-start timer is 5 seconds. After HA failover, BGP on the new primary unit has to wait for 5 seconds to connect to its neighbors, and then register BFD requests after establishing the connections. With static BFD neighbors, BFD requests and sessions can be created as soon as possible after the failover. The new command `get router info bfd requests` shows the BFD peer requests.

A BFD session created for a static BFD neighbor/peer request will initialize its state as "INIT" instead of "DOWN" and its detection time `asbfd-required-min-rx * bfd-detect-mult` milliseconds.

When a BFD control packet with nonzero `your_discr` is received, if no session can be found to match the `your_discr`, instead of discarding the packet, other fields in the packet, such as addressing information, are used to choose one session that was just initialized, with zero as its remote discriminator.

When a BFD session in the up state receives a control packet with zero as `your_discr` and down as the state, the session will change its state into down but will not notify this down event to BGP and/or other registered clients.

Link failover (port monitoring or interface monitoring)

Link failover means that if a monitored interface fails, the cluster reorganizes to reestablish a link to the network that the monitored interface was connected to and to continue operating with minimal or no disruption of network traffic.

You configure monitored interfaces (also called interface monitoring or port monitoring) by selecting the interfaces to monitor as part of the cluster HA configuration.

You can monitor up to 64 interfaces.

The interfaces that you can monitor appear on the port monitor list. You can monitor all FortiGate interfaces including redundant interfaces and 802.3ad aggregate interfaces.

You cannot monitor the following types of interfaces (you cannot select the interfaces on the port monitor list):

- FortiGate interfaces that contain an internal switch.
- VLAN subinterfaces.
- IPsec VPN interfaces.
- Individual physical interfaces that have been added to a redundant or 802.3ad aggregate interface.
- FortiGate-5000 series backplane interfaces that have not been configured as network interfaces.

If you are configuring a virtual cluster you can create a different port monitor configuration for each virtual cluster. Usually for each virtual cluster you would monitor the interfaces that have been added to the virtual domains in each virtual cluster.



Wait until after the cluster is up and running to enable interface monitoring. You do not need to configure interface monitoring to get a cluster up and running and interface monitoring will cause failovers if for some reason during initial setup a monitored interface has become disconnected. You can always enable interface monitoring once you have verified that the cluster is connected and operating properly.



You should only monitor interfaces that are connected to networks, because a failover may occur if you monitor an unconnected interface.

To enable interface monitoring - GUI

Use the following steps to monitor the port1 and port2 interfaces of a cluster.

1. Connect to the cluster GUI.
 2. Go to **System > HA** and edit the primary unit (**Role** is **MASTER**).
 3. Select the **Port Monitor** check boxes for the **port1** and **port2** interfaces and select **OK**.
- The configuration change is synchronized to all cluster units.

To enable interface monitoring - CLI

Use the following steps to monitor the port1 and port2 interfaces of a cluster.

1. Connect to the cluster CLI.
2. Enter the following command to enable interface monitoring for port1 and port2.

```
configure system ha
  set monitor port1 port2
end
```

The following example shows how to enable monitoring for the external, internal, and DMZ interfaces.

```
config system ha
  set monitor external internal dmz
end
```

With interface monitoring enabled, during cluster operation, the cluster monitors each cluster unit to determine if the monitored interfaces are operating and connected. Each cluster unit can detect a failure of its network interface hardware. Cluster units can also detect if its network interfaces are disconnected from the switch they should be connected to.



Cluster units cannot determine if the switch that its interfaces are connected to is still connected to the network. However, you can use remote IP monitoring to make sure that the cluster unit can connect to downstream network devices. See [Remote link failover on page 1556](#).

Because the primary unit receives all traffic processed by the cluster, a cluster can only process traffic from a network if the primary unit can connect to it. So, if the link between a network and the primary unit fails, to maintain communication with this network, the cluster must select a different primary unit; one that is still

connected to the network. Unless another link failure has occurred, the new primary unit will have an active link to the network and will be able to maintain communication with it.

To support link failover, each cluster unit stores link state information for all monitored cluster units in a link state database. All cluster units keep this link state database up to date by sharing link state information with the other cluster units. If one of the monitored interfaces on one of the cluster units becomes disconnected or fails, this information is immediately shared with all cluster units.

If a monitored interface on the primary unit fails

If a monitored interface on the primary unit fails, the cluster renegotiates to select a new primary unit using the process described in [Primary unit selection on page 1385](#). Because the cluster unit with the failed monitored interface has the lowest monitor priority, a different cluster unit becomes the primary unit. The new primary unit should have fewer link failures.

After the failover, the cluster resumes and maintains communication sessions in the same way as for a device failure. See [Device failover on page 1522](#).

If a monitored interface on a subordinate unit fails

If a monitored interface on a subordinate unit fails, this information is shared with all cluster units. The cluster does not renegotiate. The subordinate unit with the failed monitored interface continues to function in the cluster.

In an active-passive cluster after a subordinate unit link failover, the subordinate unit continues to function normally as a subordinate unit in the cluster.

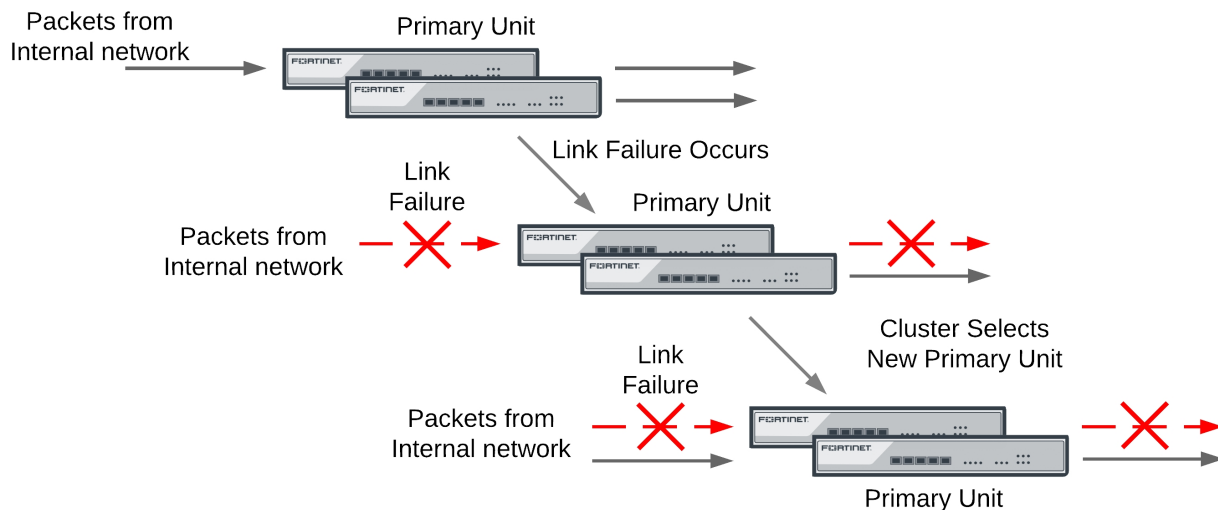
In an active-active cluster after a subordinate unit link failure:

- The subordinate unit with the failed monitored interface can continue processing connections between functioning interfaces. However, the primary unit stops sending sessions to a subordinate unit that use any failed monitored interfaces on the subordinate unit.
- If session pickup is enabled, all sessions being processed by the subordinate unit failed interface that can be failed over are failed over to other cluster units. Sessions that cannot be failed over are lost and have to be restarted.
- If session pickup is not enabled all sessions being processed by the subordinate unit failed interface are lost.

How link failover maintains traffic flow

Monitoring an interface means that the interface is connected to a high priority network. As a high priority network, the cluster should maintain traffic flow to and from the network, even if a link failure occurs. Because the primary unit receives all traffic processed by the cluster, a cluster can only process traffic from a network if the primary unit can connect to it. So, if the link that the primary unit has to a high priority network fails, to maintain traffic flow to and from this network, the cluster must select a different primary unit. This new primary unit should have an active link to the high priority network.

A link failure causes a cluster to select a new primary unit



If a monitored interface on the primary unit fails, the cluster renegotiates and selects the cluster unit with the highest monitor priority to become the new primary unit. The cluster unit with the highest monitor priority is the cluster unit with the most monitored interfaces connected to networks.

After a link failover, the primary unit processes all traffic and all subordinate units, even the cluster unit with the link failure, share session and link status. In addition all configuration changes, routes, and IPsec SAs are synchronized to the cluster unit with the link failure.

In an active-active cluster, the primary unit load balances traffic to all the units in the cluster. The cluster unit with the link failure can process connections between its functioning interfaces (for, example if the cluster has connections to an internal, external, and DMZ network, the cluster unit with the link failure can still process connections between the external and DMZ networks).

Recovery after a link failover and controlling primary unit selection (controlling falling back to the prior primary unit)

If you find and correct the problem that caused a link failure (for example, re-connect a disconnected network cable) the cluster updates its link state database and re-negotiates to select a primary unit.

What happens next depends on how the cluster configuration affects primary unit selection:

- The former primary unit will once again become the primary unit (falling back to becoming the primary unit)
- The primary unit will not change.

As described in [Primary unit selection and age on page 1387](#), when the link is restored, if no options are configured to control primary unit selection and the cluster age difference is less than 300 seconds the former

primary unit will once again become the primary unit. If the age differences are greater than 300 seconds then a new primary unit is not selected. Since you have no control on the age difference the outcome can be unpredictable. This is not a problem in cases where its not important which unit becomes the primary unit.

Preventing a primary unit change after a failed link is restored

Some organizations will not want the cluster to change primary units when the link is restored. Instead they would rather wait to restore the primary unit during a maintenance window. This functionality is not directly supported, but you can experiment with changing some primary unit selection settings. For example, in most cases it should work to enable override on all cluster units and make sure their priorities are the same. This should mean that the primary unit should not change after a failed link is restored.

Then, when you want to restore the original primary unit during a maintenance window you can just set its Device Priority higher. After it becomes the primary unit you can reset all device priorities to the same value. Alternatively during a maintenance window you could reboot the current primary unit and any subordinate units except the one that you want to become the primary unit.

If the `override` CLI keyword is enabled on one or more cluster units and the device priority of a cluster unit is set higher than the others, when the link failure is repaired and the cluster unit with the highest device priority will always become the primary unit.

Testing link failover

You can test link failure by disconnecting the network cable from a monitored interface of a cluster unit. If you disconnect a cable from a primary unit monitored interface the cluster should renegotiate and select one of the other cluster units as the primary unit. You can also verify that traffic received by the disconnected interface continues to be processed by the cluster after the failover.

If you disconnect a cable from a subordinate unit interface the cluster will not renegotiate.

Updating MAC forwarding tables when a link failover occurs

When a FortiGate HA cluster is operating and a monitored interface fails on the primary unit, the primary unit usually becomes a subordinate unit and another cluster unit becomes the primary unit. After a link failover, the new primary unit sends gratuitous ARP packets to refresh the MAC forwarding tables (also called arp tables) of the switches connected to the cluster. This is normal link failover operation.

Even when gratuitous ARP packets are sent, some switches may not be able to detect that the primary unit has become a subordinate unit and will keep sending packets to the former primary unit. This can occur if the switch does not detect the failure and does not clear its MAC forwarding table.

You have another option available to make sure the switch detects the failover and clears its MAC forwarding tables. You can use the following command to cause a cluster unit with a monitored interface link failure to briefly shut down all of its interfaces (except the heartbeat interfaces) after the failover occurs:

```
config system ha
    set link-failed-signal enable
end
```

Usually this means each interface of the former primary unit is shut down for about a second. When this happens the switch should be able to detect this failure and clear its MAC forwarding tables of the MAC addresses of the former primary unit and pickup the MAC addresses of the new primary unit. Each interface will shut down for a second but the entire process usually takes a few seconds. The more interfaces the FortiGate has, the longer it will take.

Normally, the new primary unit also sends gratuitous ARP packets that also help the switch update its MAC forwarding tables to connect to the new primary unit. If `link-failed-signal` is enabled, sending gratuitous ARP packets is optional and can be disabled if you don't need it or if its causing problems. See [Disabling gratuitous ARP packets after a failover on page 1532](#)

Multiple link failures

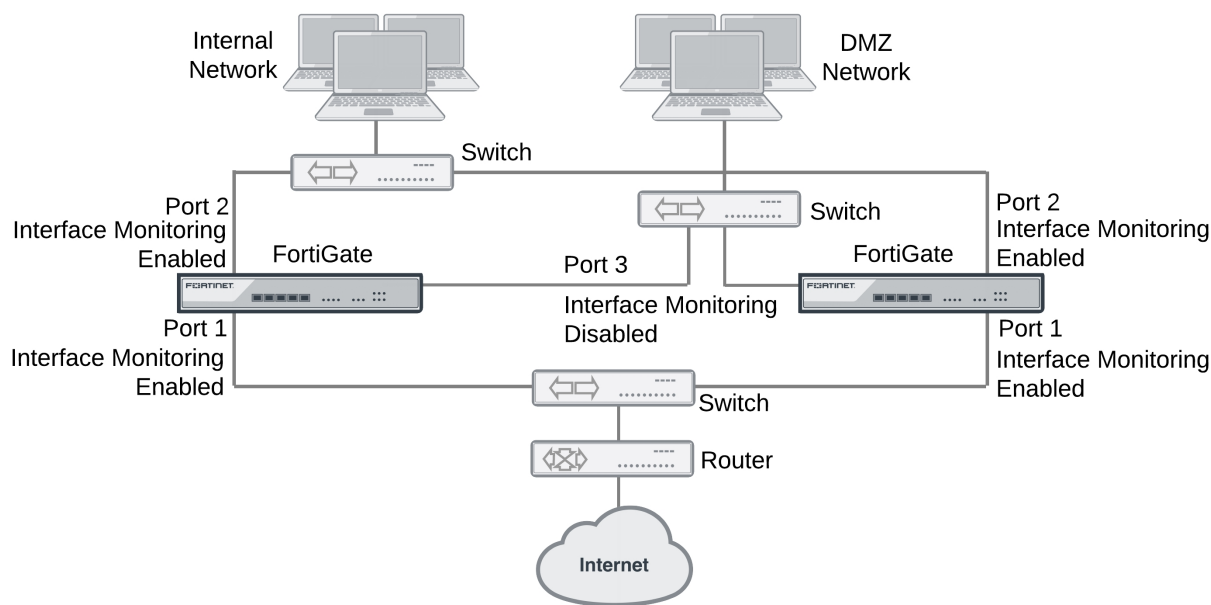
Every time a monitored interface fails, the cluster repeats the processes described above. If multiple monitored interfaces fail on more than one cluster unit, the cluster continues to negotiate to select a primary unit that can provide the most network connections.

Example link failover scenarios

For the following examples, assume a cluster configuration consisting of two FortiGates (FGT_1 and FGT_2) connected to three networks: internal using port2, external using port1, and DMZ using port3. In the HA configuration, the device priority of FGT_1 is set higher than the unit priority of FGT_2.

The cluster processes traffic flowing between the internal and external networks, between the internal and DMZ networks, and between the external and DMZ networks. If there are no link failures, FGT1 becomes the primary unit because it has the highest device priority.

Sample link failover scenario topology



Example the port1 link on FGT_1 fails

If the port1 link on FGT_1 fails, FGT_2 becomes primary unit because it has fewer interfaces with a link failure. If the cluster is operating in active-active mode, the cluster load balances traffic between the internal network

(port2) and the DMZ network (port3). Traffic between the internet (port1) and the internal network (port2) and between the internet (port1) and the DMZ network (port3) is processed by the primary unit only.

Example port2 on FGT_1 and port1 on FGT_2 fail

If port2 on FGT_1 and port1 on FGT_2 fail, then FGT_1 becomes the primary unit. After both of these link failures, both cluster units have the same monitor priority. So the cluster unit with the highest device priority (FGT_1) becomes the primary unit.

Only traffic between the internet (port1) and DMZ (port3) networks can pass through the cluster and the traffic is handled by the primary unit only. No load balancing will occur if the cluster is operating in active-active mode.

Monitoring VLAN interfaces

If the FortiGates in the cluster have VLAN interfaces, you can use the following command to monitor all VLAN interfaces and write a log message if one of the VLAN interfaces is found to be down.

Once configured, this feature works by verifying that the primary unit can connect to the subordinate unit over each VLAN. This verifies that the switch that the VLAN interfaces are connected to is configured correctly for each VLAN. If the primary unit cannot connect to the subordinate unit over one of the configured VLANs the primary unit writes a link monitor log message indicating that the named VLAN went down (log message id 20099).

Use the following CLI command to enable monitoring VLAN interfaces:

```
config system ha-monitor
  set monitor-vlan enable/disable
  set vlan-hb-interval <interval_seconds>
  set vlan-hb-lost-threshold <vlan-lost-heartbeat-threshold>
end
```

`vlan-hb-interval` is the time between sending VLAN heartbeat packets over the VLAN. The VLAN heartbeat range is 1 to 30 seconds. The default is 5 seconds.

`vlan-hb-lost-threshold` is the number of consecutive VLAN heartbeat packets that are not successfully received across the VLAN before assuming that the VLAN is down. The default value is 3, meaning that if 3 heartbeat packets sent over the VLAN are not received then the VLAN is considered to be down. The range is 1 to 60 packets.

A VLAN heartbeat interval of 5 means the time between heartbeat packets is five seconds. A VLAN heartbeat threshold of 3 means it takes $5 \times 3 = 15$ seconds to detect that a VLAN is down.

Sub-second failover

On FortiGate models 395xB and 3x40B HA link failover supports sub-second failover (that is a failover time of less than one second). Sub-second failover is available for interfaces that can issue a link failure system call when the interface goes down. When an interface experiences a link failure and sends the link failure system call, the FGCP receives the system call and initiates a link failover.

For interfaces that do not support subsection failover, port monitoring regularly polls the connection status of monitored interfaces. When a check finds that an interface has gone down, port monitoring causes a link failover. Sub-second failover results in a link failure being detected sooner because the system doesn't have to wait for the next poll to find out about the failure.

Sub-second failover can accelerate HA failover to reduce the link failover time to less than one second under ideal conditions. Actual failover performance may vary depending on traffic patterns and network configuration. For example, some network devices may respond slowly to an HA failover.

No configuration changes are required to support sub-second failover. However, for best sub-second failover results, the recommended heartbeat interval is 100ms and the recommended lost heartbeat threshold is 5 (see [Modifying heartbeat timing on page 1528](#)).

```
config system ha
    set hb-lost-threshold 5
    set hb-interval 1
end
```

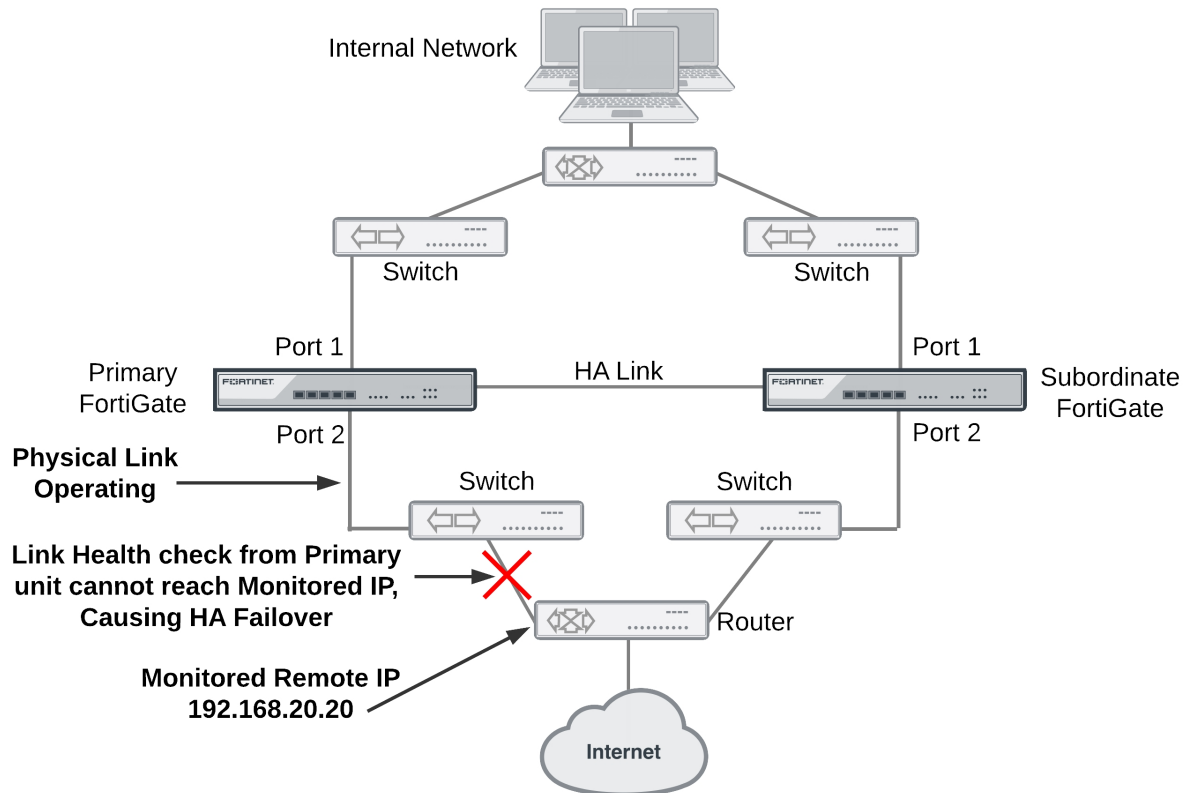
For information about how to reduce failover times, see [Failover performance on page 1566](#).

Remote link failover

Remote link failover (also called remote IP monitoring) is similar to HA port monitoring and link health monitoring (also known as dead gateway detection). Port monitoring causes a cluster to failover if a monitored primary unit interface fails or is disconnected. Remote IP monitoring uses link health monitors configured for FortiGate interfaces on the primary unit to test connectivity with IP addresses of network devices. Usually these would be IP addresses of network devices not directly connected to the cluster. For example, a downstream router. Remote IP monitoring causes a failover if one or more of these remote IP addresses does not respond to link health checking.

By being able to detect failures in network equipment not directly connected to the cluster, remote IP monitoring can be useful in a number of ways depending on your network configuration. For example, in a full mesh HA configuration, with remote IP monitoring, the cluster can detect failures in network equipment that is not directly connected to the cluster but that would interrupt traffic processed by the cluster if the equipment failed.

Example HA remote IP monitoring topology



In the simplified example topology shown above, the switch connected directly to the primary unit is operating normally but the link on the other side of the switches fails. As a result traffic can no longer flow between the primary unit and the internet.

To detect this failure you can create a link health monitor for port2 that causes the primary unit to test connectivity to 192.168.20.20. If the health monitor cannot connect to 192.268.20.20 the cluster fails over and the subordinate unit becomes the new primary unit. After the failover, the health check monitor on the new primary unit can connect to 192.168.20.20 so the failover maintains connectivity between the internal network and the internet through the cluster.

To configure remote IP monitoring

1. Enter the following commands to configure HA remote monitoring for the example topology.
 - Enter the `pingserver-monitor-interface` keyword to enable HA remote IP monitoring on port2.
 - Leave the `pingserver-failover-threshold` set to the default value of 5. This means a failover occurs if the link health monitor doesn't get a response after 5 attempts.
 - Enter the `pingserver-flip-timeout` keyword to set the flip timeout to 120 minutes. After a failover, if HA remote IP monitoring on the new primary unit also causes a failover, the flip timeout prevents the failover from occurring until the timer runs out. Setting the `pingserver-flip-timeout` to 120 means that remote IP monitoring can only cause a failover every 120 minutes. This flip timeout is required to prevent repeating failovers if remote IP monitoring causes a failover from all cluster units because none of the cluster units can connect to the monitored IP addresses.

```
config system ha
```

```

set pingserver-monitor-interface port2
set pingserver-failover-threshold 5
set pingserver-flip-timeout 120
end

```

2. Enter the following commands to add a link health monitor for the port2 interface and to set HA remote IP monitoring priority for this link health monitor.
 - Enter the `detectserver` keyword to set the health monitor server IP address to 192.168.20.20.
 - Leave the `ha-priority` keyword set to the default value of 1. You only need to change this priority if you change the HA `pingserver-failover-threshold`. The `ha-priority` setting is not synchronized among cluster units.



The `ha-priority` setting is not synchronized among cluster units. So if you want to change the `ha-priority` setting you must change it separately on each cluster unit. Otherwise it will remain set to the default value of 1.

- Use the `interval` keyword to set the time between link health checks and use the `failtime` keyword to set the number of times that a health check can fail before a failure is detected (the failover threshold). The following example reduces the failover threshold to 2 but keeps the health check interval at the default value of 5.

```

config system link-monitor
edit ha-link-monitor
set server 192.168.20.20
set srcintf port2
set ha-priority 1
set interval 5
set failtime 2
end

```

Configuring IPv6 remote IP monitoring

You can add link monitors to monitor remote IPv6 addresses. For example, use the following command:

```

config system link-monitor
edit port3
set addr-mode ipv6
set server 2001:db8:0:1
set protocol ping6
set ha-priority 1
set interval 5
set failtime 2
set gateway-ip6 2001:db8:0:2
set source-ip6 2001:db8:0:32
end

```

Adding HA remote IP monitoring to multiple interfaces

You can enable HA remote IP monitoring on multiple interfaces by adding more interface names to the `pingserver-monitor-interface` keyword. If your FortiGate configuration includes VLAN interfaces, aggregate interfaces and other interface types, you can add the names of these interfaces to the `pingserver-monitor-interface` keyword to configure HA remote IP monitoring for these interfaces.

For example, enable remote IP monitoring for interfaces named port2, port20, and vlan_234:

```

config system ha
    set pingserver-monitor-interface port2 port20 vlan_234
    set pingserver-failover-threshold 10
    set pingserver-flip-timeout 120
end

```

Then configure health monitors for each of these interfaces. In the following example, default values are accepted for all settings other than the server IP address.

```

config system link-monitor
    edit port2
        set server 192.168.20.20
    next
    edit port20
        set server 192.168.20.30
    next
    edit vlan_234
        set server 172.20.12.10
    end

```

Changing the link monitor failover threshold

If you have multiple link monitors you may want a failover to occur only if more than one of them fails.

For example, you may have 3 link monitors configured on three interfaces but only want a failover to occur if two of the link monitors fail. To do this you must set the HA priorities of the link monitors and the HA `pingserver-failover-threshold` so that the priority of one link monitor is less than the failover threshold but the added priorities of two link monitors is equal to or greater than the failover threshold. Failover occurs when the HA priority of all failed link monitors reaches or exceeds the threshold.

For example, set the failover threshold to 10 and monitor three interfaces:

```

config system ha
    set pingserver-monitor-interface port2 port20 vlan_234
    set pingserver-failover-threshold 10
    set pingserver-flip-timeout 120
end

```

Then set the HA priority of link monitor server to 5.



The HA Priority (`ha-priority`) setting is not synchronized among cluster units. In the following example, you must set the HA priority to 5 by logging into each cluster unit.

```

config system link-monitor
    edit port2
        set srcintf port2
        set server 192.168.20.20
        set ha-priority 5
    next
    edit port20
        set srcintf port20
        set server 192.168.20.30
        set ha-priority 5
    next
    edit vlan_234
        set srcintf vlan_234
        set server 172.20.12.10
    end

```

```
    set ha-priority 5
end
```

If only one of the link monitors fails, the total link monitor HA priority will be 5, which is lower than the failover threshold so a failover will not occur. If a second link monitor fails, the total link monitor HA priority of 10 will equal the failover threshold, causing a failover.

By adding multiple link monitors and setting the HA priorities for each, you can fine tune remote IP monitoring. For example, if it is more important to maintain connections to some networks you can set the HA priorities higher for these link monitors. And if it is less important to maintain connections to other networks you can set the HA priorities lower for these link monitors. You can also adjust the failover threshold so that if the cluster cannot connect to one or two high priority IP addresses a failover occurs. But a failover will not occur if the cluster cannot connect to one or two low priority IP addresses.

Monitoring multiple IP addresses from one interface

You can add multiple IP addresses to a single link monitor to use HA remote IP monitoring to monitor more than one IP address from a single interface. If you add multiple IP addresses, the health checking will be with all of the addresses at the same time. The link monitor only fails when no responses are received from all of the addresses.

```
config system link-monitor
    edit port2
        set srcintf port2
        set server 192.168.20.20 192.168.20.30 172.20.12.10
    end
```

Flip timeout

The HA remote IP monitoring configuration also involves setting a flip timeout. The flip timeout is required to reduce the frequency of failovers if, after a failover, HA remote IP monitoring on the new primary unit also causes a failover. This can happen if the new primary unit cannot connect to one or more of the monitored remote IP addresses. The result could be that until you fix the network problem that blocks connections to the remote IP addresses, the cluster will experience repeated failovers. You can control how often the failovers occur by setting the flip timeout. The flip timeout stops HA remote IP monitoring from causing a failover until the primary unit has been operating for the duration of the flip timeout.

If you set the flip timeout to a relatively high number of minutes you can find and repair the network problem that prevented the cluster from connecting to the remote IP address without the cluster experiencing very many failovers. Even if it takes a while to detect the problem, repeated failovers at relatively long time intervals do not usually disrupt network traffic.

Use the following command to set the flip timeout to 3 hours (360 minutes):

```
config system ha
    set pingserver-flip-timeout 360
end
```

Restoring normal cluster operation after the remote link is restored

In a remote IP monitoring configuration, if you also want the same cluster unit to always be the primary unit you can set its device priority higher and enable override. With this configuration, when a remote IP monitoring failover occurs, after the flip timeout expires another failover will occur (because override is enabled) and the unit with override enabled becomes the primary unit again. So the cluster automatically returns to normal operation.

The primary unit starts remote IP monitoring again. If the remote link is restored the cluster continues to operate normally. If, however, the remote link is still down, remote link failover causes the cluster to failover again. This will repeat each time the flip timeout expires until the failed remote link is restored.

You can use the `pingserver-slave-force-reset` option to control this behavior. By default this option is enabled and the behavior described above occurs. The overall behavior is that when the remote link is restored the cluster automatically returns to normal operation after the flip timeout.

If you disable `pingserver-slave-force-reset` after the initial remote IP monitoring failover nothing will happen after the flip timeout (as long as the new primary unit doesn't experience some kind of failover). The result is that repeated failovers no longer happen. But it also means that the original primary unit will remain the subordinate unit and will not resume operating as the primary unit.

Detecting HA remote IP monitoring failovers

Just as with any HA failover, you can detect HA remote IP monitoring failovers by using SNMP to monitor for HA traps. You can also use alert email to receive notifications of HA status changes and monitor log messages for HA failover log messages. In addition, FortiGates send the critical log message `Ping Server is down` when a ping server fails. The log message includes the name of the interface that the ping server has been added to.

Failover and attached network equipment

It normally takes a cluster approximately 6 seconds to complete a failover. However, the actual failover time experienced by your network users may depend on how quickly the switches connected to the cluster interfaces accept the cluster MAC address update from the primary unit. If the switches do not recognize and accept the gratuitous ARP packets and update their MAC forwarding table, the failover time will increase.

Also, individual session failover depends on whether the cluster is operating in active-active or active-passive mode, and whether the content of the traffic is to be virus scanned. Depending on application behavior, it may take a TCP session a longer period of time (up to 30 seconds) to recover completely.

Monitoring cluster units for failover

You can use logging and SNMP to monitor cluster units for failover. Both the primary and subordinate units can be configured to write log messages and send SNMP traps if a failover occurs. You can also log into the cluster GUI and CLI to determine if a failover has occurred.

NAT/Route mode active-passive cluster packet flow

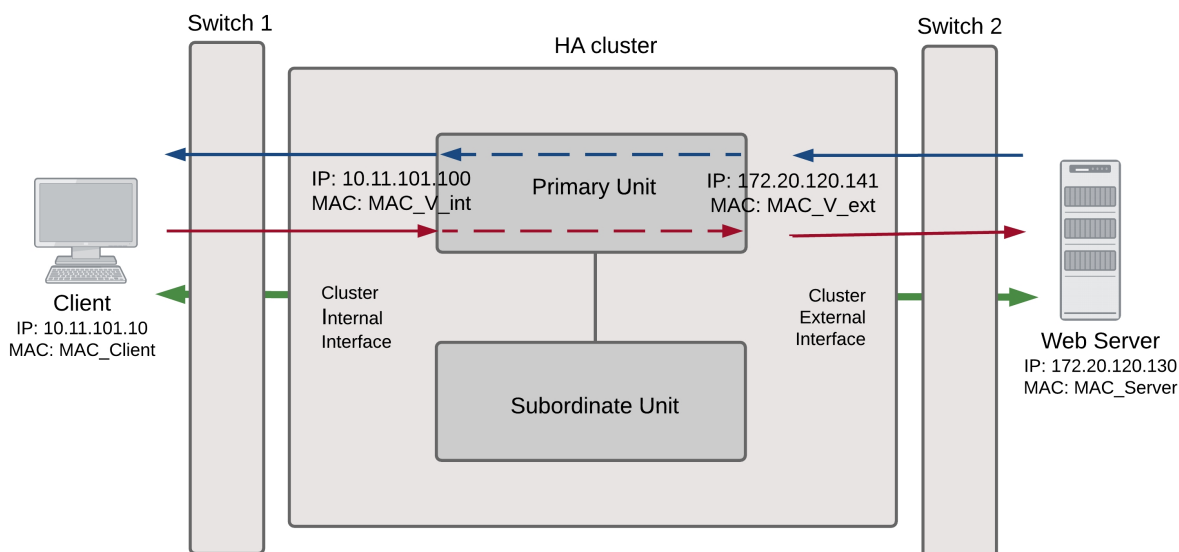
This section describes how packets are processed and how failover occurs in an active-passive HA cluster running in NAT/Route mode. In the example, the NAT/Route mode cluster acts as the internet firewall for a client computer's internal network. The client computer's default route points at the IP address of the cluster internal interface. The client connects to a web server on the internet. internet routing routes packets from the cluster external interface to the web server, and from the web server to the cluster external interface.

In an active-passive cluster operating in NAT/Route mode, four MAC addresses are involved in communication between the client and the web server when the primary unit processes the connection:

- Internal virtual MAC address (MAC_V_int) assigned to the primary unit internal interface,
- External virtual MAC address (MAC_V_ext) assigned to the primary unit external interface,
- Client MAC address (MAC_Client),
- Server MAC address (MAC_Server),

In NAT/Route mode, the HA cluster works as a gateway when it responds to ARP requests. Therefore, the client and server only know the gateway MAC addresses. The client only knows the cluster internal virtual MAC address (MAC_V_int) and the server only know the cluster external virtual MAC address (MAC_V_ext).

NAT/Route mode active-passive packet flow



Packet flow from client to web server

1. The client computer requests a connection from 10.11.101.10 to 172.20.120.130.
2. The default route on the client computer recognizes 10.11.101.100 (the cluster IP address) as the gateway to the external network where the web server is located.
3. The client computer issues an ARP request to 10.11.101.100.
4. The primary unit intercepts the ARP request, and responds with the internal virtual MAC address (MAC_V_int) which corresponds to its IP address of 10.11.101.100.
5. The client's request packet reaches the primary unit internal interface.

	IP address	MAC address
Source	10.11.101.10	MAC_Client
Destination	172.20.120.130	MAC_V_int

6. The primary unit processes the packet.
7. The primary unit forwards the packet from its external interface to the web server.

	IP address	MAC address
Source	172.20.120.141	MAC_V_ext
Destination	172.20.120.130	MAC_Server

8. The primary unit continues to process packets in this way unless a failover occurs.

Packet flow from web server to client

1. When the web server responds to the client's packet, the cluster external interface IP address (172.20.120.141) is recognized as the gateway to the internal network.
2. The web server issues an ARP request to 172.20.120.141.
3. The primary unit intercepts the ARP request, and responds with the external virtual MAC address (MAC_V_ext) which corresponds its IP address of 172.20.120.141.
4. The web server then sends response packets to the primary unit external interface.

	IP address	MAC address
Source	172.20.120.130	MAC_Server
Destination	172.20.120.141	MAC_V_ext

5. The primary unit processes the packet.
6. The primary unit forwards the packet from its internal interface to the client.

	IP address	MAC address
Source	172.20.120.130	MAC_V_int
Destination	10.11.101.10	MAC_Client

7. The primary unit continues to process packets in this way unless a failover occurs.

When a failover occurs

The following steps are followed after a device or link failure of the primary unit causes a failover.

1. If the primary unit fails the subordinate unit becomes the primary unit.
2. The new primary unit changes the MAC addresses of all of its interfaces to the HA virtual MAC addresses.
The new primary unit has the same IP addresses and MAC addresses as the failed primary unit.
3. The new primary unit sends gratuitous ARP packets from the internal interface to the 10.11.101.0 network to associate its internal IP address with the internal virtual MAC address.
4. The new primary unit sends gratuitous ARP packets to the 172.20.120.0 to associate its external IP address with the external virtual MAC address.
5. Traffic sent to the cluster is now received and processed by the new primary unit.
If there were more than two cluster units in the original cluster, these remaining units would become subordinate units.

Transparent mode active-passive cluster packet flow

This section describes how packets are processed and how failover occurs in an active-passive HA cluster running in transparent mode. The cluster is installed on an internal network in front of a mail server and the client connects to the mail server through the transparent mode cluster.

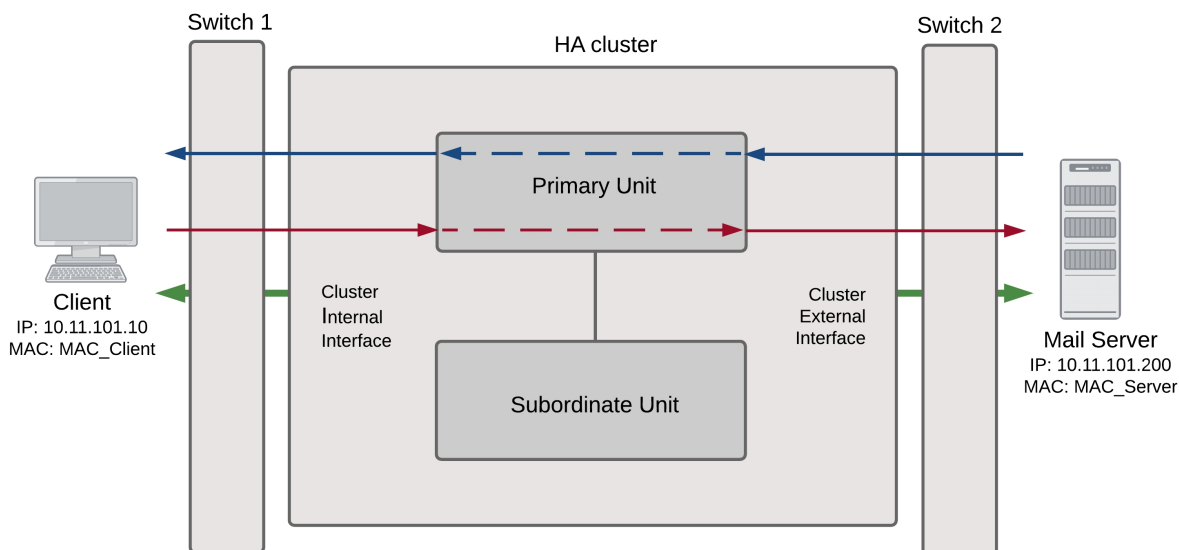
In an active-passive cluster operating in transparent mode, two MAC addresses are involved in the communication between a client and a server when the primary unit processes a connection:

- Client MAC address (MAC_Client)
- Server MAC address (MAC_Server)

The HA virtual MAC addresses are not directly involved in communication between the client and the server. The client computer sends packets to the mail server and the mail server sends responses. In both cases the packets are intercepted and processed by the cluster.

The cluster's presence on the network is transparent to the client and server computers. The primary unit sends gratuitous ARP packets to Switch 1 that associate all MAC addresses on the network segment connected to the cluster external interface with the HA virtual MAC address. The primary unit also sends gratuitous ARP packets to Switch 2 that associate all MAC addresses on the network segment connected to the cluster internal interface with the HA virtual MAC address. In both cases, this results in the switches sending packets to the primary unit interfaces.

Transparent mode active-passive packet flow



Packet flow from client to mail server

1. The client computer requests a connection from 10.11.101.10 to 10.11.101.200.
2. The client computer issues an ARP request to 10.11.101.200.
3. The primary unit forwards the ARP request to the mail server.
4. The mail server responds with its MAC address (MAC_Server) which corresponds to its IP address of 10.11.101.200. The primary unit returns the ARP response to the client computer.
5. The client's request packet reaches the primary unit internal interface.

IP address	MAC address
------------	-------------

Source	10.11.101.10	MAC_Client
Destination	10.11.101.200	MAC_Server

- The primary unit processes the packet.
- The primary unit forwards the packet from its external interface to the mail server.

	IP address	MAC address
Source	10.11.101.10	MAC_Client
Destination	10.11.101.200	MAC_Server

- The primary unit continues to process packets in this way unless a failover occurs.

Packet flow from mail server to client

- To respond to the client computer, the mail server issues an ARP request to 10.11.101.10.
- The primary unit forwards the ARP request to the client computer.
- The client computer responds with its MAC address (MAC_Client) which corresponds to its IP address of 10.11.101.10. The primary unit returns the ARP response to the mail server.
- The mail server's response packet reaches the primary unit external interface.

	IP address	MAC address
Source	10.11.101.200	MAC_Server
Destination	10.11.101.10	MAC_Client

- The primary unit processes the packet.
- The primary unit forwards the packet from its internal interface to the client.

	IP address	MAC address
Source	10.11.101.200	MAC_Server
Destination	10.11.101.10	MAC_Client

- The primary unit continues to process packets in this way unless a failover occurs.

When a failover occurs

The following steps are followed after a device or link failure of the primary unit causes a failover.

- If the primary unit fails, the subordinate unit negotiates to become the primary unit.
- The new primary unit changes the MAC addresses of all of its interfaces to the HA virtual MAC address.
- The new primary unit sends gratuitous ARP packets to switch 1 to associate its MAC address with the MAC addresses on the network segment connected to the external interface.
- The new primary unit sends gratuitous ARP packets to switch 2 to associate its MAC address with the MAC addresses on the network segment connected to the internal interface.
- Traffic sent to the cluster is now received and processed by the new primary unit.

If there were more than two cluster units in the original cluster, these remaining units would become subordinate units.

Failover performance

This section describes the designed device and link failover times for a FortiGate cluster and also shows results of a failover performance test.

Device failover performance

By design FGCP device failover time is 2 seconds for a two-member cluster with ideal network and traffic conditions. If sub-second failover is enabled the failover time can drop below 1 second.

All cluster units regularly receive HA heartbeat packets from all other cluster units over the HA heartbeat link. If any cluster unit does not receive a heartbeat packet from any other cluster unit for 2 seconds, the cluster unit that has not sent heartbeat packets is considered to have failed.

It may take another few seconds for the cluster to negotiate and re-distribute communication sessions. Typically if sub-second failover is not enabled you can expect a failover time of 9 to 15 seconds depending on the cluster and network configuration. The failover time can also be increased by more complex configurations and or configurations with network equipment that is slow to respond.

You can change the `hb-lost-threshold` to increase or decrease the device failover time. See [Modifying heartbeat timing on page 1528](#) for information about using `hb-lost-threshold`, and other heartbeat timing settings.

Link failover performance

Link failover time is controlled by how long it takes for a cluster to synchronize the cluster link database. When a link failure occurs, the cluster unit that experienced the link failure uses HA heartbeat packets to broadcast the updated link database to all cluster units. When all cluster units have received the updated database the failover is complete.

It may take another few seconds for the cluster to negotiate and re-distribute communication sessions.

Reducing failover times

- Keep the network configuration as simple as possible with as few as possible network connections to the cluster.
- If possible operate the cluster in transparent mode.
- Use high-performance switches to that the switches failover to interfaces connected to the new primary unit as quickly as possible.
- Use accelerated FortiGate interfaces. In some cases accelerated interfaces will reduce failover times.
- Make sure the FortiGate sends multiple gratuitous arp packets after a failover. In some cases, sending more gratuitous arp packets will cause connected network equipment to recognize the failover sooner. To send 10 gratuitous arp packets:

```
config system ha
  set arps 10
end
```

- Reduce the time between gratuitous arp packets. This may also caused connected network equipment to recognize the failover sooner. To send 50 gratuitous arp packets with 1 second between each packet:

```
config system ha
  set arps 50
```

```
    set arps-interval 1
end
```

- Reduce the number of lost heartbeat packets and reduce the heartbeat interval timers to be able to more quickly detect a device failure. To set the lost heartbeat threshold to 3 packets and the heartbeat interval to 100 milliseconds:

```
config system ha
    set hb-interval 1
    set hb-lost-threshold 3
end
```

- Reduce the hello state hold down time to reduce the amount of the time the cluster waits before transitioning from the hello to the work state. To set the hello state hold down time to 5 seconds:

```
config system ha
    set hello-holddown 5
end
```

- Enable sending a link failed signal after a link failover to make sure that attached network equipment responds as quickly as possible to a link failure. To enable the link failed signal:

```
config system ha
    set link-failed-signal enable
end
```

Session failover (session-pickup)

Session failover means that after the primary unit fails, communications sessions resume on the new primary unit with minimal or no interruption. Two categories of sessions need to be resumed after a failover:

- Sessions passing through the cluster
- Sessions terminated by the cluster

If you enable session failover (also called session-pickup) for the cluster, during cluster operation the primary unit informs the subordinate units of changes to the primary unit connection and state tables for sessions passing through the cluster, keeping the subordinate units up-to-date with the traffic currently being processed by the cluster. All synchronization activity takes place over the HA heartbeat link using TCP/703 and UDP/703 packets.

After a failover the new primary unit recognizes open sessions that were being handled by the cluster. The sessions continue to be processed by the new primary unit and are handled according to their last known state.



Session-pickup has some limitations. For example, session failover is not supported for sessions being scanned by proxy-based security profiles. Session failover is supported for sessions being scanned by flow-based security profiles; however, flow-based sessions that fail over are not inspected after they fail over. For more limitations, see [Session failover limitations for sessions passing through the cluster on page 1570](#).

Sessions terminated by the cluster include management sessions (such as HTTPS connections to the FortiGate GUI or SSH connection to the CLI as well as SNMP and logging and so on). Also included in this category are IPsec VPN, SSL VPN, sessions terminated by the cluster, explicit proxy, WAN Optimization and web caching. In general, whether or not session-pickup is enabled, these sessions do not failover and have to be restarted. There are some exceptions though, particularly for IPsec and SSL VPN. For more information, see [Session failover limitations for sessions terminated by the cluster on page 1573](#).

Enabling session-pickup for TCP, UDP, ICMP, and multicast session failover

To enable session-pickup, go to **System > HA** and enable session-pickup.

From the CLI enter:

```
config system ha
    set session-pickup enable
end
```

When session-pickup is enabled, the FGCP synchronizes the primary unit's TCP session table to all cluster units. As soon as a new TCP session is added to the primary unit session table, that session is synchronized to all cluster units. This synchronization happens as quickly as possible to keep the session tables synchronized.

If the primary unit fails, the new primary unit uses its synchronized session table to resume all TCP sessions that were being processed by the former primary unit with only minimal interruption. Under ideal conditions all TCP sessions should be resumed. This is not guaranteed though and under less than ideal conditions some TCP sessions may need to be restarted.

Enabling UDP and ICMP session failover

If session pickup is enabled, you can use the following command to also enable UDP and ICMP session failover:

```
config system ha
  set session-pickup-connectionless enable
end
```

Enabling multicast session failover

To configure multicast session failover, use the following command to change the multicast TTL timer to a smaller value than the default. The recommended setting to support multicast session failover is 120 seconds (2 minutes). The default setting is 600 seconds (10 minutes).

```
config system ha
  set multicast-ttl 120
end
```

The multicast TTL timer controls how long to keep synchronized multicast routes on the backup unit (so they are present on the backup unit when it becomes the new primary unit after a failover). If you set the multicast TTL lower the multicast routes on the backup unit are refreshed more often so are more likely to be accurate. Reducing this time causes route synchronization to happen more often and could affect performance.

If session pickup is disabled

If you leave session pickup disabled, the cluster does not keep track of sessions and after a failover, active sessions have to be restarted or resumed. Most session can be resumed as a normal result of how TCP/IP communications resumes communication after any routine network interruption.



The session-pickup setting does not affect session failover for sessions terminated by the cluster.

If you do not require session failover protection, leaving session pickup disabled may reduce CPU usage and reduce HA heartbeat network bandwidth usage. Also if your cluster is mainly being used for traffic that is not synchronized (for example, for proxy-based security profile processing) enabling session pickup is not recommended since most sessions will not be failed over anyway.

If session pickup is not enabled, the FGCP does not synchronize the primary unit session table to other cluster units and sessions do not resume after a failover. After a device or link failover all sessions are briefly interrupted and must be re-established at the application level after the cluster renegotiates.

Many protocols can successfully restart sessions with little, if any, loss of data. For example, after a failover, users browsing the web can just refresh their browsers to resume browsing. Since most HTTP sessions are very short, in most cases they will not even notice an interruption unless they are downloading large files. Users downloading a large file may have to restart their download after a failover.

Other protocols may experience data loss and some protocols may require sessions to be manually restarted. For example, a user downloading files with FTP may have to either restart downloads or restart their FTP client.

Improving session synchronization performance

Two HA configuration options are available to reduce the performance impact of enabling session-pickup. They include reducing the number of sessions that are synchronized by adding a session pickup delay and using more FortiGate interfaces for session synchronization.

Reducing the number of sessions that are synchronized

If session pickup is enabled, as soon as new sessions are added to the primary unit session table they are synchronized to the other cluster units. Enable the `session-pickup-delay` CLI option to reduce the number of sessions that are synchronized by synchronizing sessions only if they remain active for more than 30 seconds. Enabling this option could greatly reduce the number of sessions that are synchronized if a cluster typically processes very many short duration sessions, which is typical of most HTTP traffic for example.

Use the following command to enable a 30 second session pickup delay:

```
config system ha
    set session-pickup-delay enable
end
```

Enabling session pickup delay means that if a failover occurs more sessions may not be resumed after a failover. In most cases short duration sessions can be restarted with only a minor traffic interruption. However, if you notice too many sessions not resuming after a failover you might want to disable this setting.

Using multiple FortiGate interfaces for session synchronization

Using the `session-sync-dev` option you can select one or more FortiGate interfaces to use for synchronizing sessions as required for session pickup. Normally session synchronization occurs over the HA heartbeat link. Using this HA option means only the selected interfaces are used for session synchronization and not the HA heartbeat link. If you select more than one interface, session synchronization traffic is load balanced among the selected interfaces.

Moving session synchronization from the HA heartbeat interface reduces the bandwidth required for HA heartbeat traffic and may improve the efficiency and performance of the cluster, especially if the cluster is synchronizing a large number of sessions. Load balancing session synchronization among multiple interfaces can further improve performance and efficiency if the cluster is synchronizing a large number of sessions.

Use the following command to perform cluster session synchronization using the port10 and port12 interfaces.

```
config system ha
    set session-sync-dev port10 port12
end
```

Session synchronization packets use Ethertype 0x8892. The interfaces to use for session synchronization must be connected together either directly using the appropriate cable (possible if there are only two units in the cluster) or using switches. If one of the interfaces becomes disconnected the cluster uses the remaining interfaces for session synchronization. If all of the session synchronization interfaces become disconnected, session synchronization reverts back to using the HA heartbeat link. All session synchronization traffic is between the primary unit and each subordinate unit.

Since large amounts of session synchronization traffic can increase network congestion, it is recommended that you keep this traffic off of your network by using dedicated connections for it.

Session failover limitations for sessions passing through the cluster

This section contains information about session failover for communication sessions passing through the cluster. In general, if session pickup is enabled, session failover is supported for most TCP traffic. This section describes details about how this all works.

Protocol	Session Failover?
Most TCP sessions.	Supported if session-pickup is enabled. (More about TCP session failover on page 1571)
Multicast sessions	Supported if multicast session-pickup is enabled. (Enabling multicast session failover).
IPv6, NAT64, and NAT66	Supported if session-pickup is enabled.
Proxy-based security profile sessions	Not Supported, sessions have to be restarted. Proxy-based features require the FortiGate to maintain very large amounts of internal state information for each session. The FGCP does not synchronize this internal state information. As a result, proxy-based sessions are not failed over. Active-active clusters can resume some of these sessions after a failover. (Active-active HA subordinate units sessions can resume after a failover on page 1573)
Flow-based security profile sessions.	Supported if session-pickup is enabled. Flow-based sessions failover, but internal state information is not synchronized so sessions that fail over are no longer inspected by security profile functions. If both flow-based and proxy-based security profile features are applied to a TCP session, that session will not resume after a failover.
UDP and ICMP, or broadcast sessions	Supported if connectionless session-pickup is enabled. Otherwise, sessions have to be restarted. (UDP, ICMP, and broadcast packet session failover on page 1572)
GPRS Tunneling Protocol (GTP)	Supported with limitations. (FortiOS Carrier GTP session failover on page 1573)
SIP	Supported for active-passive HA only. (SIP session failover on page 1572)
SIMPLE, or SCCP signal session	Not supported, sessions have to be restarted.
SSL offloading and HTTP multiplexing	Not supported, sessions have to be restarted. (SSL offloading and HTTP multiplexing session failover on page 1573)

More about TCP session failover

TCP sessions that are not being processed by security profile features resume after a failover even if these sessions are accepted by security policies with security profiles. Only TCP sessions that are actually being processed by these security profile features do not resume after a failover. For example:

- TCP sessions that are not virus scanned, web filtered, spam filtered, content archived, or are not SIP, SIMPLE, or SCCP signal traffic resume after a failover, even if they are accepted by a security policy with security profile options enabled. For example, SNMP TCP sessions through the FortiGate resume after a failover because FortiOS does not apply any security profile options to SNMP sessions.
- TCP sessions for a protocol for which security profile features have not been enabled resume after a failover even if they are accepted by a security policy with security profile features enabled. For example, if you have not enabled any antivirus or content archiving settings for FTP, FTP sessions resume after a failover.

UDP, ICMP, and broadcast packet session failover

By default, even with session pickup enabled, the FGCP does not maintain a session table for UDP, ICMP, or broadcast packets. So the cluster does not specifically support failover of these packets.

Some UDP traffic can continue to flow through the cluster after a failover. This can happen if, after the failover, a UDP packet that is part of an already established communication stream matches a security policy. Then a new session will be created and traffic will flow. So after a short interruption, UDP sessions can appear to have failed over. However, this may not be reliable for the following reasons:

- UDP packets in the direction of the security policy must be received before reply packets can be accepted. For example, if a port1 -> port2 policy accepts UDP packets, UDP packets received at port2 destined for the network connected to port1 will not be accepted until the policy accepts UDP packets at port1 that are destined for the network connected to port2. So, if a user connects from an internal network to the internet and starts receiving UDP packets from the internet (for example streaming media), after a failover the user will not receive any more UDP packets until the user re-connects to the internet site.
- UDP sessions accepted by NAT policies will not resume after a failover because NAT will usually give the new session a different source port. So only traffic for UDP protocols that can handle the source port changing during a session will continue to flow.

You can however, enable session pickup for UDP and ICMP packets by enabling session pickup for TCP sessions and then enabling session pickup for connectionless sessions:

```
config system ha
    set session-pickup enable
    set session-pickup-connectionless enable
end
```

This configuration causes the cluster units to synchronize UDP and ICMP session tables and if a failover occurs UDP and ICMP sessions are maintained.

SIP session failover

If session pickup is enabled, the FGCP supports SIP session failover (also called stateful failover) for active-passive HA.

SIP session failover replicates SIP states to all cluster units. If an HA failover occurs, all in-progress SIP calls (setup complete) and their RTP flows are maintained and the calls will continue after the failover with minimal or no interruption.

SIP calls being set up at the time of a failover may lose signaling messages. In most cases the SIP clients and servers should use message retransmission to complete the call setup after the failover has completed. As a result, SIP users may experience a delay if their calls are being set up when an HA failover occurs. But in most cases the call setup should be able to continue after the failover.

FortiOS Carrier GTP session failover

FortiOS Carrier HA supports GTP session failover. The primary unit synchronizes the GTP tunnel state to all cluster units after the GTP tunnel setup is completed. After the tunnel setup is completed, GTP sessions use UDP and HA does not synchronize UDP sessions to all cluster units. However, similar to other UDP sessions, after a failover, since the new primary unit will have the GTP tunnel state information, GTP UDP sessions using the same tunnel can continue to flow with some limitations.

The limitation on packets continuing to flow is that there has to be a security policy to accept the packets. For example, if the FortiOS Carrier unit has an internal to external security policy, GTP UDP sessions using an established tunnel that are received by the internal interface are accepted by the security policy and can continue to flow. However, GTP UDP packets for an established tunnel that are received at the external interface cannot flow until packets from the same tunnel are received at the internal interface.

If you have bi-directional policies that accept GTP UDP sessions then traffic in either direction that uses an established tunnel can continue to flow after a failover without interruption.

SSL offloading and HTTP multiplexing session failover

SSL offloading and HTTP multiplexing are both enabled from firewall virtual IPs and firewall load balancing. Similar to the features applied by security profile, SSL offloading and HTTP multiplexing requires the FortiGate to maintain very large amounts of internal state information for each session. Sessions accepted by security policies containing virtual IPs or virtual servers with SSL offloading or HTTP multiplexing enabled do not resume after a failover.

Active-active HA subordinate units sessions can resume after a failover

In an active-active cluster, subordinate units process sessions. After a failover, all cluster units that are still operating may be able to continue processing the sessions that they were processing before the failover. These sessions are maintained because after the failover the new primary unit uses the HA session table to continue to send session packets to the cluster units that were processing the sessions before the failover. Cluster units maintain their own information about the sessions that they are processing and this information is not affected by the failover. In this way, the cluster units that are still operating can continue processing their own sessions without loss of data.

The cluster keeps processing as many sessions as it can. But some sessions can be lost. Depending on what caused the failover, sessions can be lost in the following ways:

- A cluster unit fails (the primary unit or a subordinate unit). All sessions that were being processed by that cluster unit are lost.
- A link failure occurs. All sessions that were being processed through the network interface that failed are lost.

This mechanism for continuing sessions is not the same as session failover because:

- Only the sessions that can be maintained are maintained.
- The sessions are maintained on the same cluster units and not re-distributed.
- Sessions that cannot be maintained are lost.

Session failover limitations for sessions terminated by the cluster

This section contains information about session failover for communication sessions terminated by the cluster. Sessions terminated by the cluster include management sessions as well as IPsec and SSL VPN, WAN Optimization and so on between the cluster and a client.

In general, most sessions terminated by the cluster have to be restarted after a failover. There are some exceptions though, for example, the FGCP provides failover for IPsec and SSL VPN sessions terminated by the cluster.



The session pickup setting does not affect session failover for sessions terminated by the cluster. Also other cluster settings such as active-active or active-passive mode do not affect session failover for sessions terminated by the cluster.

Protocol	Session Failover?
Administrative or management connections such as connecting to the GUI or CLI, SNMP, syslog, communication with FortiManager, FortiAnalyzer and so on.	Not supported, sessions have to be restarted.
Explicit web proxy, WCCP, WAN Optimization and Web Caching.	Not supported, sessions have to be restarted. (Explicit web proxy, explicit FTP proxy, WCCP, WAN optimization and Web Caching session failover on page 1574)
IPsec VPN tunnels terminating at the FortiGate	Supported. SAs and related IPsec VPN tunnel data is synchronized to cluster members. (Synchronizing IPsec VPN SAs on page 1575)
SSL VPN tunnels terminating at the FortiGate	Partially supported. Sessions are not synchronized and have to be restarted. Authentication failover and cookie failover is supported. Once the client restarts the session they shouldn't have to re-authenticate. (SSL VPN session failover and SSL VPN authentication failover on page 1575)
PPTP and L2TP VPN terminating at the FortiGate	Not supported, sessions have to be restarted.

Explicit web proxy, explicit FTP proxy, WCCP, WAN optimization and Web Caching session failover

Explicit web proxy, explicit FTP proxy, WCCP, WAN optimization and web caching sessions all require the FortiGate to maintain very large amounts of internal state information for each session. This information is not maintained and these sessions do not resume after a failover.

The active-passive HA clustering is recommended for WAN optimization. All WAN optimization sessions are processed by the primary unit only. Even if the cluster is operating in active-active mode, HA does not load-balance WAN optimization sessions.

Also, Web cache and byte cache databases are only stored on the primary unit. These databases are not synchronized to the cluster. So, after a failover, the new primary unit must rebuild its web and byte caches. As well, the new primary unit cannot connect to a SAS partition that the failed primary unit used.

Rebuilding the byte caches can happen relatively quickly because the new primary unit gets byte cache data from the other FortiGates that it is participating with in WAN optimization tunnels.

SSL VPN session failover and SSL VPN authentication failover

Session failover is not supported for SSL VPN tunnels. However, authentication failover is supported for the communication between the SSL VPN client and the FortiGate. This means that after a failover, SSL VPN clients can re-establish the SSL VPN session between the SSL VPN client and the FortiGate without having to authenticate again.

However, all sessions inside the SSL VPN tunnel that were running before the failover are stopped and have to be restarted. For example, file transfers that were in progress would have to be restarted. As well, any communication sessions with resources behind the FortiGate that are started by an SSL VPN session have to be restarted.

To support SSL VPN cookie failover, when an SSL VPN session starts, the FGCP distributes the cookie created to identify the SSL VPN session to all cluster units.

PPTP and L2TP VPN sessions

PPTP and L2TP VPNs are supported in HA mode. For a cluster you can configure PPTP and L2TP settings and you can also add security policies to allow PPTP and L2TP pass through. However, the FGCP does not provide session failover for PPTP or L2TP. After a failover, all active PPTP and L2TP sessions are lost and must be restarted.

Synchronizing IPsec VPN SAs

The FGCP synchronizes IPsec security associations (SAs) between cluster members so that if a failover occurs, the cluster can resume IPsec sessions without having to establish new SAs. The result is improved failover performance because IPsec sessions are not interrupted to establish new SAs. Also, establishing a large number of SAs can reduce cluster performance.

The FGCP implements slightly different synchronization mechanisms for IKEv1 and IKEv2.

Synchronizing SAs for IKEv1

When an SA is synchronized to the subordinate units, the sequence number is set to the maximum sequence number. After a failover, all inbound traffic that connects with the new primary unit and uses the SA will be accepted without needing to re-key. However, first outbound packet to use the SA causes the sequence number to overflow and so causes the new primary unit to re-key the SA.

Please note the following:

- The cluster synchronizes all IPsec SAs.
- IPsec SAs are not synchronized until the IKE process has finished synchronizing the ISAKMP SAs. This is required in for dialup tunnels since it is the synchronizing of the ISAKMP SA that creates the dialup tunnel.
- A dialup interface is created as soon as the phase 1 is complete. This ensures that the when HA synchronizes phase 1 information the dialup name is included.
- If the IKE process re-starts for any reason it deletes any dialup tunnels that exist. This forces the peer to re-key them.
- IPsec SA deletion happens immediately. Routes associated with a dialup tunnel that is being deleted are cleaned up synchronously as part of the delete, rather than waiting for the SA hard-expiry.
- The FGCP does not sync the IPsec tunnel MTU from the primary unit to the subordinate units. This means that after HA failover if the first packet received by the FortiGate arrives after the HA route has been deleted and before the new route is added and the packet is larger than the default MTU of 1024 then the FortiGate sends back an

ICMP fragmentation required. However, as soon as routing is re-established then the MTU will be corrected and traffic will flow.

Synchronizing SAs for IKEv2

Due to the way the IKEv2 protocol is designed the FGCP cannot use exactly the same solution that is used for synchronizing IKEv1 SAs, though it is similar.

For IKEv2, like IKEv1, the FGCP synchronizes IKE and ISAKMP SAs from the primary unit to the subordinate units. However, for IKEv2 the FGCP cannot actually use this IKE SA to send/receive IKE traffic because IKEv2 includes a sequence number in every IKE message and thus it would require synchronizing every message to the subordinate units to keep the sequence numbers on the subordinate units up to date.

Instead, the FGCP synchronizes IKEv2 Message IDs. This Message ID Sync allows IKEv2 to re-negotiate send and receive message ID counters after a failover. By doing this, the established IKE SA can remain up, instead of re-negotiating.

The `diagnose vpn ike stats` command shows statistics for the number of HA messages sent/received for IKEv2. The output of this command includes a number of fields prefixed with `ha` that contain high availability related-data. For example:

```
.
.
.
ha.resync: 0
ha.vike.sync: 0
ha.conn.sync: 0
ha.sync.tx: 1
ha.sync.rx: 0
ha.sync.rx.len.bad: 0
.
.
.
```

WAN optimization and HA

You can configure WAN optimization on a FortiGate HA cluster. The recommended HA configuration for WAN optimization is active-passive mode. Also, when the cluster is operating, all WAN optimization sessions are processed by the primary unit only. Even if the cluster is operating in active-active mode, HA does not load-balance WAN optimization sessions. HA also does not support WAN optimization session failover.

In a cluster, the primary unit only stores web cache and byte cache databases. These databases are not synchronized to the subordinate units. So, after a failover, the new primary unit must rebuild its web and byte caches. As well, the new primary unit cannot connect to a SAS partition that the failed primary unit used.

Rebuilding the byte caches can happen relatively quickly because the new primary unit gets byte cache data from the other FortiGates that it is participating with in WAN optimization tunnels.

HA and load balancing

FGCP active-active (a-a) load balancing distributes network traffic among all of the units in a cluster. Load balancing can improve cluster performance because the processing load is shared among multiple cluster units.

This chapter describes how active-active load balancing works and provides detailed active-active HA NAT/Route and transparent mode packet flow descriptions.

Load balancing overview

FGCP active-active HA uses a technique similar to unicast load balancing in which the primary unit is associated with the cluster HA virtual MAC addresses and cluster IP addresses. The primary unit is the only cluster unit to receive packets sent to the cluster. The primary unit then uses a load balancing schedule to distribute sessions to all of the units in the cluster (including the primary unit). Subordinate unit interfaces retain their actual MAC addresses and the primary unit communicates with the subordinate units using these MAC addresses. Packets exiting the subordinate units proceed directly to their destination and do not pass through the primary unit first.

By default, active-active HA load balancing distributes proxy-based security profile processing to all cluster units. Proxy-based security profile processing is CPU and memory-intensive, so FGCP load balancing may result in higher throughput because resource-intensive processing is distributed among all cluster units.

Proxy-based security profile processing that is load balanced includes proxy-based virus scanning, proxy-based web filtering, proxy-based email filtering, and proxy-based data leak prevention (DLP) of HTTP, FTP, IMAP, IMAPS, POP3, POP3S, SMTP, SMTPS, IM, and NNTP, sessions accepted by security policies.

Other features enabled in security policies such as Endpoint security, traffic shaping and authentication have no effect on active-active load balancing.

You can also enable `load-balance-all` to have the primary unit load balance all TCP sessions. Load balancing TCP sessions increases overhead and may actually reduce performance so it is disabled by default. You can also enable `load-balance-udp` to have the primary unit load balance all UDP sessions. Load balancing UDP sessions also increases overhead so it is also disabled by default.

NP4 and NP6 processors can also offload and accelerate load balancing.

During active-active HA load balancing the primary unit uses the configured load balancing schedule to determine the cluster unit that will process a session. The primary unit stores the load balancing information for each load balanced session in the cluster load balancing session table. Using the information in this table, the primary unit can then forward all of the remaining packets in each session to the appropriate cluster unit. The load balancing session table is synchronized among all cluster units.

HTTPS, ICMP, multicast, and broadcast sessions are never load balanced and are always processed by the primary unit. IPS, Application Control, flow-based virus scanning, flow-based web filtering, flow-based DLP, flow-based email filtering, VoIP, IM, P2P, IPsec VPN, HTTPS, SSL VPN, HTTP multiplexing, SSL offloading, WAN optimization, explicit web proxy, and WCCP sessions are also always processed only by the primary unit.

In addition to load balancing, active-active HA also provides the same session, device and link failover protection as active-passive HA. If the primary unit fails, a subordinate unit becomes the primary unit and resumes operating the cluster.

Active-active HA also maintains as many load balanced sessions as possible after a failover by continuing to process the load balanced sessions that were being processed by the cluster units that are still operating. See [Active-active HA subordinate units sessions can resume after a failover on page 1](#) for more information.

Load balancing schedules

The load balancing schedule controls how the primary unit distributes packets to all cluster units. You can select from the following load balancing schedules.

Schedule	Description
None	No load balancing. Select None when the cluster interfaces are connected to load balancing switches. If you select None, the Primary unit does not load balance traffic and the subordinate units process incoming traffic that does not come from the Primary unit. For all other load balancing schedules, all traffic is received first by the Primary unit, and then forwarded to the subordinate units. The subordinate units only receive and process packets sent from the primary unit.
Hub	Load balancing if the cluster interfaces are connected to a hub. Traffic is distributed to cluster units based on the source IP and destination IP of the packet.
Least-Connection	If the cluster units are connected using switches, select Least Connection to distribute network traffic to the cluster unit currently processing the fewest connections.
Round-Robin	If the cluster units are connected using switches, select Round-Robin to distribute network traffic to the next available cluster unit.
Weighted Round-Robin	Similar to round robin, but weighted values are assigned to each of the units in a cluster based on their capacity and on how many connections they are currently processing. For example, the primary unit should have a lower weighted value because it handles scheduling and forwards traffic. Weighted round robin distributes traffic more evenly because units that are not processing traffic will be more likely to receive new connections than units that are very busy.
Random	If the cluster units are connected using switches, select Random to randomly distribute traffic to cluster units.
IP	Load balancing according to IP address. If the cluster units are connected using switches, select IP to distribute traffic to units in a cluster based on the source IP and destination IP of the packet.
IP Port	Load balancing according to IP address and port. If the cluster units are connected using switches, select IP Port to distribute traffic to units in a cluster based on the source IP, source port, destination IP, and destination port of the packet.

Once a packet has been propagated to a subordinate unit, all packets are part of that same communication session are also propagated to that same subordinate unit. Traffic is distributed according to communication session, not just according to individual packet.

Any subordinate unit that receives a forwarded packet processes it, without applying load balancing. Note that subordinate units are still considered to be active, because they perform routing, virus scanning, and other

FortiGate tasks on their share of the traffic. Active subordinate units also share their session and link status information with all cluster units. The only things that active members do not do is make load balancing decisions.

Even though the primary unit is responsible for the load balancing process, the primary unit still acts like a FortiGate in that it processes packets, performing, routing, firewall, virus scanning, and other FortiGate tasks on its share of the traffic. Depending on the load balancing schedule used, the primary unit may assign itself a smaller share of the total load.

More about active-active failover

If a subordinate unit fails, the primary unit re-distributes the sessions that the subordinate was processing among the remaining active cluster members. If the primary unit fails, the subordinate units negotiate to select a new primary unit. The new primary unit continues to distribute packets among the remaining active cluster units.

Failover works in a similar way if the cluster consists of only two units. If the primary unit fails the subordinate unit negotiates and becomes the new primary unit. If the subordinate unit fails, the primary unit processes all traffic. In both cases, the single remaining unit continues to function as a primary unit, maintaining the HA virtual MAC address for all of its interfaces.

HTTPS sessions, active-active load balancing, and proxy servers

To prevent HTTPS web filtering problems active-active HA does not load balance HTTPS sessions. The FortiGate identifies HTTPS sessions as all sessions received on the HTTPS TCP port. The default HTTPS port is 443. You can go to **Policy & Objects > Policy > SSL/SSH Inspection** to use a custom port for HTTPS sessions. If you change the HTTPS port, the FGCP stops load balancing all sessions that use the custom HTTPS port.

Normally you would not change the HTTPS port. However, if your network uses a proxy server for HTTPS traffic you may have to change to the custom HTTPS port used by your proxy server. If your network uses a proxy server you might also use the same port for both HTTP and HTTPS traffic. In this case you would configure the FortiGate to use custom ports for both HTTP and HTTPS traffic. Go to **Policy & Objects > Policy > Proxy Options** to use a custom port for HTTP.

Using the same port for HTTP and HTTPS traffic can cause problems with active-active clusters because active-active clusters always load balance HTTP traffic. If both HTTP and HTTPS use the same port, the active-active cluster cannot differentiate between HTTP and HTTPS traffic and will load balance both.

As mentioned above, load balancing HTTPS traffic may cause problems with HTTPS web filtering. To avoid this problem, you should configure your proxy server to use different ports for HTTP and HTTPS traffic. Then configure your cluster to also use different ports for HTTP and HTTPS.

Selecting a load balancing schedule

You can select the load balancing schedule when initially configuring the cluster and you can change the load balancing schedule at any time while the cluster is operating without affecting cluster operation.

You can select a load balancing schedule from the CLI. Use the following command to select a load balancing schedule:

```
config system ha
  set schedule {hub | ip | ipport | leastconnection | none | random | round-robin
    | weight-round-robin}
end
```

Load balancing TCP and UDP sessions

You can use the following command to configure the cluster to load balance TCP sessions in addition to security profile sessions.

```
config system ha
    set load-balance-all enable
end
```

Enabling `load-balance-all` to add load balancing of TCP sessions may not improve performance because the cluster requires additional overhead to load balance sessions. Load balancing a TCP session usually requires about as much overhead as just processing it. On the other hand, TCP load balancing performance may be improved if your FortiGate includes NP4 or NP6 processors.

You can enable `load-balance-all` and monitor network performance to see if it improves. If performance is not improved, you might want to change the HA mode to active-passive since active-active HA is not providing any benefit.

On some FortiGate models you can use the following command to also load balance UDP sessions:

```
config system ha
    set load-balance-udp enable
end
```

Similar to load balancing TCP sessions, load balancing UDP sessions may also not improve performance. Also UDP load balancing performance may be improved with NP4 and NP6 processors.

Using NP4 or NP6 processors to offload load balancing

FortiGates that include NP4 and NP6 network processors can provide hardware acceleration for active-active HA cluster by offloading load balancing from the primary unit CPU. Network processors are especially useful when load balancing TCP and UDP sessions.

The first packet of every new session is received by the primary unit and the primary unit uses its load balancing schedule to select the cluster unit that will process the new session. This information is passed back to the network processor and all subsequent packets of the same sessions are offloaded to the network processor which sends the packet directly to a subordinate unit. Load balancing is effectively offloaded from the primary unit to the network processor resulting in a faster and more stable active-active cluster.

To take advantage of network processor load balancing acceleration, connect the cluster unit interfaces with network processors to the busiest networks. Connect other interfaces to less busy networks. No special FortiOS or HA configuration is required. Network processor acceleration of active-active HA load balancing is supported for any active-active HA configuration or active-active HA load balancing schedule.

Configuring weighted-round-robin weights

You can configure weighted round-robin load balancing for a cluster and configure the static weights for each of the cluster units according to their priority in the cluster. When you set `schedule` to `weight-round-robin` you can use the `weight` option to set the static weight of each cluster unit. The static weight is set according to the priority of each unit in the cluster. A FortiGate HA cluster can contain up to four FortiGates so you can set up to 4 static weights.

The priority of a cluster unit is determined by its device priority, the number of monitored interfaces that are functioning, its age in the cluster and its serial number. Priorities are used to select a primary unit and to set the priorities of all of the subordinate units. Thus the priority of a cluster unit can change depending on configuration

settings, link failures and so on. Since weights are also set using this priority, the weights are independent of specific cluster units but do depend on the role of the each unit in the cluster.

You can use the following command to display the relative priorities of the units in a cluster. The cluster unit serial numbers and their priorities are listed in the last few lines of the command output. This example shows a cluster of three FortiGates:

```
get system ha status
.
.
.
Slave : FG-5KD3914800284, operating cluster index = 1
Master: FG-5KD3914800344, operating cluster index = 0
Slave : FG-5KD3914800353, operating cluster index = 2
```

The primary unit always has the highest priority and the subordinate units have lower priorities.

The default static weight for each cluster unit is 40. This means that sessions are distributed evenly among all cluster units. You can use the `set weight` command to change the static weights of cluster units to distribute sessions to cluster units depending on their priority in the cluster. The weight can be between 0 and 255. Increase the weight to increase the number of connections processed by the cluster unit with that priority.

You set the weight for each unit separately. For the example cluster of 3 FortiGates you can set the weight for each unit as follows:

```
config system ha
  set mode a-a
  set schedule weight-round-robin
  set weight 0 5
  set weight 1 10
  set weight 2 15
end
```

If you enter the `get` command to view the HA configuration the output for `weight` would be:

```
weight 5 10 15 40 40 40 40 40 40 40 40 40 40 40 40
```

This configuration has the following results if the output of the `get system ha status` command is that shown above:

- The first five connections are processed by the primary unit (priority 0, weight 5).
- The next 10 connections are processed by the first subordinate unit (priority 1, weight 10)
- The next 15 connections are processed by the second subordinate unit (priority 2, weight 15)

Dynamically optimizing weighted load balancing according to how busy cluster units are

In conjunction with using static weights to load balance sessions among cluster units you can configure a cluster to dynamically load balance sessions according to individual cluster unit CPU usage, memory usage, and number of HTTP, FTP, IMAP, POP3, SMTP, or NNTP proxy-based security profile sessions. If any of these system loading indicators increases above configured thresholds, weighted load balancing dynamically sends fewer new sessions to the busy unit until it recovers.

High CPU or memory usage indicates that a unit is under increased load and may not be able to process more sessions. HTTP, FTP, IMAP, POP3, SMTP, or NNTP proxy use are also good indicators of how busy a cluster unit is, since processing high numbers of these proxy sessions can quickly reduce overall cluster unit performance.

For example, you can set a CPU usage high watermark threshold. When a cluster unit reaches this high watermark threshold fewer sessions are sent to it. With fewer sessions to process the cluster unit's CPU usage should fall back to the low watermark threshold. When the low watermark threshold is reached the cluster resumes normal load balancing of sessions to the cluster unit.

You can set individual high and low watermark thresholds and weights for CPU usage, memory usage, and for the number of HTTP, FTP, IMAP, POP3, SMTP, or NNTP proxy sessions.

The CPU usage, memory usage, and proxy weights determine how the cluster load balances sessions when a high watermark threshold is reached and also affect how the cluster load balances sessions when multiple cluster units reach different high watermark thresholds at the same time. For example, you might be less concerned about a cluster unit reaching the memory usage high watermark threshold than reaching the CPU usage high watermark threshold. If this is the case you can set the weight lower for memory usage. Then, if one cluster unit reaches the CPU usage high watermark threshold and a second cluster unit reaches the memory usage high watermark threshold the cluster will load balance more sessions to the cluster unit with high memory usage and fewer sessions to the cluster unit with high CPU usage. As a result, reaching the CPU usage high watermark will have a greater affect on how sessions are redistributed than reaching the memory usage high watermark.

When a high watermark threshold is reached, the corresponding weight is subtracted from the static weight of the cluster unit. The lower the weight the fewer the number of sessions that are load balanced to that unit.

Subsequently when the low watermark threshold is reached, the static weight of the cluster unit returns to its configured value. For the weights to all be effective the weights assigned to the load indicators should usually be lower than or equal to the static weights assigned to the cluster units.

Use the following command to set thresholds and weights for CPU and memory usage and HTTP, FTP, IMAP, POP3, SMTP, or NNTP proxy sessions:

```
config system ha
  set mode a-a
  set schedule weight-round-robin
  set cpu-threshold <weight> <low> <high>
  set memory-threshold <weight> <low> <high>
  set http-proxy-threshold <weight> <low> <high>
  set ftp-proxy-threshold <weight> <low> <high>
  set imap-proxy-threshold <weight> <low> <high>
  set nntp-proxy-threshold <weight> <low> <high>
  set pop3-proxy-threshold <weight> <low> <high>
  set smtp-proxy-threshold <weight> <low> <high>
end
```

For each option, the weight range is 0 to 255 and the default weight is 5. The low and high watermarks are a percent (0 to 100). The default low and high watermarks are 0 which means they are disabled. The default configuration when weighted load balancing is enabled looks like the following:

```
config system ha
  set mode a-a
  set schedule weight-round-robin
  set cpu-threshold 5 0 0
  set memory-threshold 5 0 0
  set http-proxy-threshold 5 0 0
  set ftp-proxy-threshold 5 0 0
  set imap-proxy-threshold 5 0 0
  set nntp-proxy-threshold 5 0 0
  set pop3-proxy-threshold 5 0 0
  set smtp-proxy-threshold 5 0 0
end
```



When you first enable HA weighted load balancing, the weighted load balancing configuration is synchronized to all cluster units and each cluster unit has the default configuration shown above. Changes to the CPU, memory, HTTP, FTP, IMAP, NNTP, POP3, and SMTP proxy thresholds and low and high watermarks must be made for each cluster unit and are not synchronized to the other cluster units.

When you configure them, the high watermarks must be greater than their corresponding low watermarks.

For CPU and memory usage the low and high watermarks are compared with the percentage CPU and memory use of the cluster unit. For each of the proxies the high and low watermarks are compared to a number that represents percent of the max number of proxy sessions being used by a proxy. This number is calculated using the formula:

$$\text{proxy usage} = (\text{current sessions} * 100) / \text{max sessions}$$

where:

`current sessions` is the number of active sessions for the proxy type.

`max sessions` is the session limit for the proxy type. The session limit depends on the FortiGate and its configuration.

You can use the following command to display the maximum and current number of sessions for a proxy:

```
get test {ftpd | http | imap | nntp | pop3 | smtp} 4
```

You can use the following command to display the maximum number of sessions and the current number of sessions for all of the proxies:

```
get test proxyworker 4
```

The command output includes lines similar to the following:

```
get test http 4
HTTP Common
Current Connections          5000/8032
```

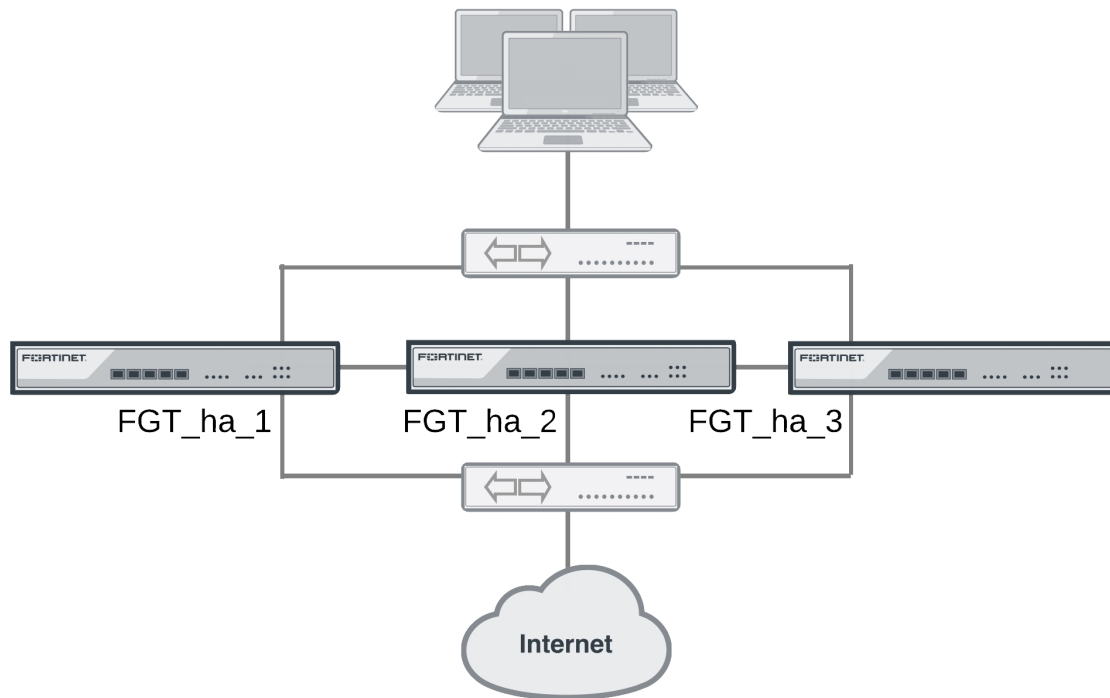
In the example, 5000 is the current number of proxy connections being used by HTTP and 8032 is the maximum number of proxy sessions allowed. For this example the proxy usage would be:

```
proxy usage = (5000 * 100) / 8032
proxy usage = 62%
```

Example weighted load balancing configuration

Consider a cluster of three FortiGates with host names FGT_ha_1, FGT_ha_2, and FGT_ha_3 as shown below. This example describes how to configure weighted load balancing settings for CPU and memory usage for the cluster and then to configure HTTP and POP3 proxy weights to send most HTTP and POP3 proxy sessions to different cluster units.

Example HA weighted load balancing configuration



Connect to the cluster CLI and use the following command to set the CPU usage threshold weight to 30, low watermark to 60, and high watermark to 80. This command also sets the memory usage threshold weight to 10, low watermark to 60, and high watermark to 90.

```

config system ha
    set mode a-a
    set schedule weight-round-robin
    set cpu-threshold 30 60 80
    set memory-threshold 10 60 90
end
  
```

The static weights for the cluster units remain at the default values of 40. Since this command changes the mode to `a-a` and the schedule to `weight-round-robin` for the first time, the weight settings are synchronized to all cluster units.

As a result of this configuration, if the CPU usage of any cluster unit (for example, FGT_ha_1) reaches 80% the static weight for that cluster unit is reduced from 40 to 10 and only 10 of every 120 new sessions are load balanced to this cluster unit. If the memory usage of FGT_ha_1 also reaches 90% the static weight further reduces to 0 and no new sessions are load balanced to FGT_ha_1. Also, if the memory usage of FGT_ha_2 reaches 90% the static weight of FGT_ha_2 reduces to 30 and 30 of every 120 new sessions are load balanced to FGT_ha_2.

Now that you have established the weight load balancing configuration for the entire cluster you can monitor the cluster to verify that processing gets distributed evenly to all cluster units. From the GUI you can go to **System > HA > View HA Statistics** and see the CPU usage, active sessions, memory usage and other statistics for all of the cluster units. If you notice that one cluster unit is more or less busy than others you can adjust the dynamic weights separately for each cluster unit.

For example, in some active-active clusters the primary unit may tend to be busier than other cluster units because in addition to processing sessions the primary unit also receives all packets sent to the cluster and performs load balancing to distribute the sessions to other cluster units. To reduce the load on the primary unit you could reduce the CPU and memory usage high watermark thresholds for the primary unit so that fewer sessions are distributed to the primary unit. You could also reduce the primary unit's high watermark setting for the proxies to distribute more proxy sessions to other cluster units.



This would only be useful if you are using device priorities and override settings to make sure the same unit always becomes the primary unit. See [Controlling primary unit selection using device priority and override on page 1395](#).

If the example cluster is configured for FGT_ha_2 to be the primary unit, connect to the FGT_ha_2's CLI and enter the following command to set CPU usage, memory usage, and proxy usage high watermark thresholds lower.

```
config system ha
  set cpu-threshold 30 60 70
  set memory-threshold 30 60 70
  set http-proxy-threshold 30 60 70
  set ftp-proxy-threshold 30 60 70
  set imap-proxy-threshold 30 60 70
  set nntp-proxy-threshold 30 60 70
  set pop3-proxy-threshold 30 60 70
  set smtp-proxy-threshold 30 60 70
end
```

As a result, when any of these factors reaches 70% on the primary unit, fewer sessions will be processed by the primary unit, preventing the number of sessions being processed from rising.

NAT/Route mode active-active cluster packet flow


This section describes an example of how packets are load balanced and how failover occurs in an active-active HA cluster running in NAT/Route mode. In the example, the NAT/Route mode cluster acts as the internet firewall for a client computer's internal network. The client computer's default route points at the IP address of the cluster internal interface. The client connects to a web server on the internet. Internet routing routes packets from the cluster external interface to the web server, and from the web server to the cluster external interface.

In NAT/Route mode, eight MAC addresses are involved in active-active communication between the client and the web server when the primary unit load balances packets to the subordinate unit:

- Internal virtual MAC address (MAC_V_int) assigned to the primary unit internal interface,
- External virtual MAC address (MAC_V_ext) assigned to the primary unit external interface,
- Client MAC address (MAC_Client),
- Server MAC address (MAC_Server),
- Primary unit original internal MAC address (MAC_P_int),
- Primary unit original external MAC address (MAC_P_ext),
- Subordinate unit internal MAC address (MAC_S_int),
- Subordinate unit external MAC address (MAC_S_ext).

In NAT/Route mode, the HA cluster works as a gateway when it responds to ARP requests. Therefore, the client and server only know the gateway MAC addresses. The client only knows the cluster internal virtual MAC address (MAC_V_int) and the server only knows the cluster external virtual MAC address (MAC_V_ext).

NAT/Route mode active-active packet flow



Packet flow from client to web server

1. The client computer requests a connection from 10.11.101.10 to 172.20.120.130.
2. The default route on the client computer recognizes 10.11.101.100 (the cluster IP address) as the gateway to the external network where the web server is located.
3. The client computer issues an ARP request to 10.11.101.100.
4. The primary unit intercepts the ARP request, and responds with the internal virtual MAC address (MAC_V_int) which corresponds to its IP address of 10.11.101.100.
5. The client's request packet reaches the primary unit internal interface.

	IP address	MAC address
Source	10.11.101.10	MAC_Client
Destination	172.20.120.130	MAC_V_int

6. The primary unit decides that the subordinate unit should handle this packet, and forwards it to the subordinate unit internal interface. The source MAC address of the forwarded packet is changed to the actual MAC address of the primary unit internal interface.

IP address	MAC address
------------	-------------

Source	10.11.101.10	MAC_P_int
Destination	172.20.120.130	MAC_S_int

7. The subordinate unit recognizes that the packet has been forwarded from the primary unit and processes it.
8. The subordinate unit forwards the packet from its external interface to the web server.

	IP address	MAC address
Source	172.20.120.141	MAC_S_ext
Destination	172.20.120.130	MAC_Server

9. The primary unit forwards further packets in the same session to the subordinate unit.
10. Packets for other sessions are load balanced by the primary unit and either sent to the subordinate unit or processed by the primary unit.

Packet flow from web server to client

1. When the web server responds to the client's packet, the cluster external interface IP address (172.20.120.141) is recognized as the gateway to the internal network.
2. The web server issues an ARP request to 172.20.120.141.
3. The primary unit intercepts the ARP request, and responds with the external virtual MAC address (MAC_V_ext) which corresponds its IP address of 172.20.120.141.
4. The web server then sends response packets to the primary unit external interface.

	IP address	MAC address
Source	172.20.120.130	MAC_Server
Destination	172.20.120.141	MAC_V_ext

5. The primary unit decides that the subordinate unit should handle this packet, and forwards it to the subordinate unit external interface. The source MAC address of the forwarded packet is changed to the actual MAC address of the primary unit external interface.

	IP address	MAC address
Source	172.20.120.130	MAC_P_ext
Destination	172.20.120.141	MAC_S_ext

6. The subordinate unit recognizes that packet has been forwarded from the primary unit and processes it.
7. The subordinate unit forwards the packet from its internal interface to the client.

	IP address	MAC address
Source	172.20.120.130	MAC_S_int
Destination	10.11.101.10	MAC_Client

8. The primary unit forwards further packets in the same session to the subordinate unit.
9. Packets for other sessions are load balanced by the primary unit and either sent to the subordinate unit or processed by the primary unit.

When a failover occurs

The following steps are followed after a device or link failure of the primary unit causes a failover.

1. If the primary unit fails, the subordinate unit negotiates to become the primary unit.
2. The new primary unit changes the MAC addresses of all of its interfaces to the HA virtual MAC addresses.
The new primary unit has the same IP addresses and MAC addresses as the failed primary unit.
3. The new primary unit sends gratuitous ARP packets to the 10.10.101.0 network to associate its internal IP address with the internal virtual MAC address.
4. The new primary unit sends gratuitous ARP packets to the 172.20.120.0 network to associate its external IP address with the external virtual MAC address.
5. Traffic sent to the cluster is now received and processed by the new primary unit.

If there were more than two cluster units in the original cluster, the new primary unit would load balance packets to the remaining cluster members.

Transparent mode active-active cluster packet flow

This section describes an example of how packets are load balanced and how failover occurs in an active-active HA cluster running in transparent mode. The cluster is installed on an internal network in front of a mail server and the client connects to the mail server through the transparent mode cluster.

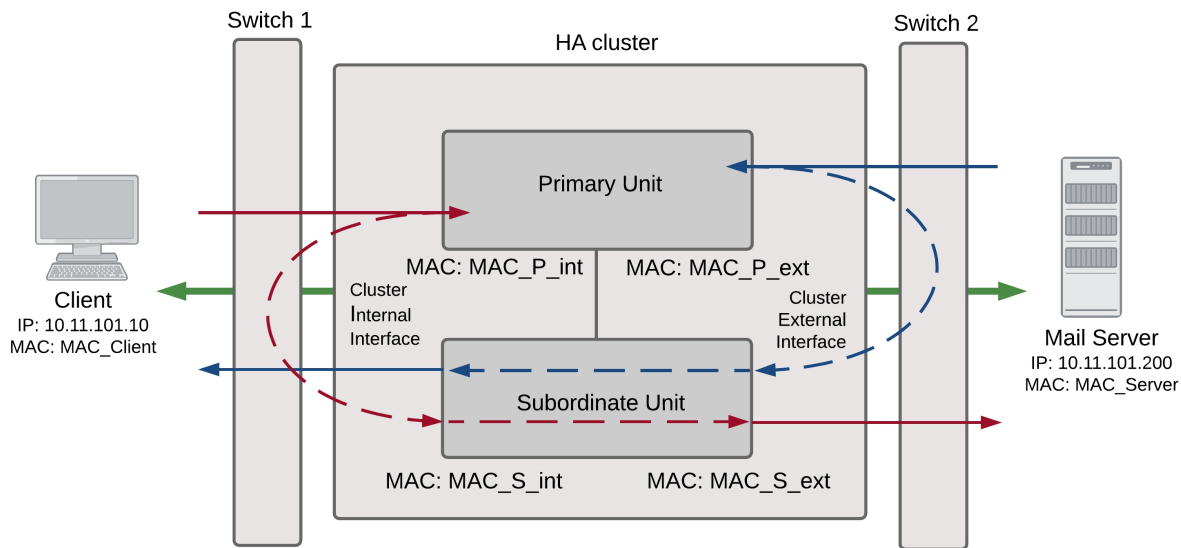
In transparent mode, six MAC addresses are involved in active-active communication between a client and a server when the primary unit load balances packets to the subordinate unit:

- Client MAC address (MAC_Client),
- Server MAC address (MAC_Server),
- Primary unit original internal MAC address (MAC_P_int),
- Primary unit original external MAC address (MAC_P_ext),
- Subordinate unit internal MAC address (MAC_S_int),
- Subordinate unit external MAC address (MAC_S_ext).

The HA virtual MAC addresses are not directly involved in communication between the client and the server. The client computer sends packets to the mail server and the mail server sends responses. In both cases the packets are intercepted and load balanced among cluster members.

The cluster's presence on the network and its load balancing are transparent to the client and server computers. The primary unit sends gratuitous ARP packets to Switch 1 that associate all MAC addresses on the network segment connected to the cluster external interface with the external virtual MAC address. The primary unit also sends gratuitous ARP packets to Switch 2 that associate all MAC addresses on the network segment connected to the cluster internal interface with the internal virtual MAC address. In both cases, this results in the switches sending packets to the primary unit interfaces.

Transparent mode active-active packet flow



Packet flow from client to mail server

1. The client computer requests a connection from 10.11.101.10 to 10.11.101.200.
2. The client computer issues an ARP request to 10.11.101.200.
3. The primary unit forwards the ARP request to the mail server.
4. The mail server responds with its MAC address (MAC_Server) which corresponds to its IP address of 10.11.101.200. The primary unit returns the ARP response to the client computer.
5. The client's request packet reaches the primary unit internal interface.

	IP address	MAC address
Source	10.11.101.10	MAC_Client
Destination	10.11.101.200	MAC_Server

6. The primary unit decides that the subordinate unit should handle this packet, and forwards it to the subordinate unit internal interface. The source MAC address of the forwarded packet is changed to the actual MAC address of the primary unit internal interface.

	IP address	MAC address
Source	10.11.101.10	MAC_P_int
Destination	10.11.101.200	MAC_S_int

7. The subordinate unit recognizes that packet has been forwarded from the primary unit and processes it.
8. The subordinate unit forwards the packet from its external interface to the mail server.

	IP address	MAC address
Source	10.11.101.10	MAC_S_ext
Destination	10.11.101.200	MAC_Server

9. The primary unit forwards further packets in the same session to the subordinate unit.
10. Packets for other sessions are load balanced by the primary unit and either sent to the subordinate unit or processed by the primary unit.

Packet flow from mail server to client

1. To respond to the client computer, the mail server issues an ARP request to 10.11.101.10.
2. The primary unit forwards the ARP request to the client computer.
3. The client computer responds with its MAC address (MAC_Client) which corresponds to its IP address of 10.11.101.10. The primary unit returns the ARP response to the mail server.
4. The mail server's response packet reaches the primary unit external interface.

	IP address	MAC address
Source	10.11.101.200	MAC_Server
Destination	10.11.101.10	MAC_Client

5. The primary unit decides that the subordinate unit should handle this packet, and forwards it to the subordinate unit external interface. The source MAC address of the forwarded packet is changed to the actual MAC address of the primary unit external interface.

	IP address	MAC address
Source	10.11.101.200	MAC_P_ext
Destination	10.11.101.10	MAC_S_ext

6. The subordinate unit recognizes that packet has been forwarded from the primary unit and processes it.
7. The subordinate unit forwards the packet from its internal interface to the client.

	IP address	MAC address
Source	10.11.101.200	MAC_S_int
Destination	10.11.101.10	MAC_Client

8. The primary unit forwards further packets in the same session to the subordinate unit.
9. Packets for other sessions are load balanced by the primary unit and either sent to the subordinate unit or processed by the primary unit.

When a failover occurs

The following steps are followed after a device or link failure of the primary unit causes a failover.

1. If the primary unit fails the subordinate unit negotiates to become the primary unit.
 2. The new primary unit changes the MAC addresses of all of its interfaces to the HA virtual MAC address.
 3. The new primary unit sends gratuitous ARP requests to switch 1 to associate its MAC address with the MAC addresses on the network segment connected to the external interface.
 4. The new primary unit sends gratuitous ARP requests to switch 2 to associate its MAC address with the MAC addresses on the network segment connected to the internal interface.
 5. Traffic sent to the cluster is now received and processed by the new primary unit.
- If there were more than two cluster units in the original cluster, the new primary unit would load balance packets to the remaining cluster members.

HA with FortiGate-VM and third-party products

This chapter provides information about operating FortiOS VM cluster and operating FortiGate clusters with third party products such as layer-2 and layer-3 switches.

FortiGate-VM for VMware HA configuration

If you want to combine two or more FortiGate-VM instances into a FortiGate Clustering Protocol (FGCP) High Availability (HA) cluster the VMware server's virtual switches used to connect the heartbeat interfaces must operate in promiscuous mode. This permits HA heartbeat communication between the heartbeat interfaces. HA heartbeat packets are non-TCP packets that use Ethertype values 0x8890, 0x8891, and 0x8893. All synchronization activity takes place over the HA heartbeat link using TCP/703 and UDP/703 packets. The FGCP uses link-local IPv4 addresses in the 169.254.0.x range for HA heartbeat interface IP addresses.

To enable promiscuous mode in VMware:

1. In the vSphere client, select your VMware server in the left pane and then select the **Configuration** tab in the right pane.
2. In **Hardware**, select **Networking**.
3. Select **Properties** of a virtual switch used to connect heartbeat interfaces.
4. In the **Properties** window left pane, select **vSwitch** and then select **Edit**.
5. Select the **Security** tab, set **Promiscuous Mode** to **Accept**, then select **OK**.
6. Select **Close**.

You must also set the virtual switches connected to other FortiGate interfaces to allow MAC address changes and to accept forged transmits. This is required because the FGCP sets virtual MAC addresses for all FortiGate interfaces and the same interfaces on the different VM instances in the cluster will have the same virtual MAC addresses.

To make the required changes in VMware:

1. In the vSphere client, select your VMware server in the left pane and then select the **Configuration** tab in the right pane.
2. In **Hardware**, select **Networking**.
3. Select **Properties** of a virtual switch used to connect FortiGate VM interfaces.
4. Set **MAC Address Changes** to **Accept**.
5. Set **Forged Transmits** to **Accept**.

FortiGate-VM for Hyper-V HA configuration

Promiscuous mode and support for MAC address spoofing is required for FortiGate-VM for Hyper-V to support FortiGate Clustering Protocol (FGCP) high availability (HA). By default the FortiGate-VM for Hyper-V has promiscuous mode enabled in the XML configuration file in the FortiGate-VM Hyper-V image. If you have problems with HA mode, confirm that this is still enabled.

In addition, because the FGCP applies virtual MAC addresses to FortiGate data interfaces and because these virtual MAC addresses mean that matching interfaces of different FortiGate-VM instances will have the same virtual MAC addresses you have to configure Hyper-V to allow MAC spoofing. But you should only enable MAC

spoofing for FortiGate-VM data interfaces. You should not enable MAC spoofing for FortiGate HA heartbeat interfaces.

With promiscuous mode enabled and the correct MAC spoofing settings you should be able to configure HA between two or more FortiGate-VM for Hyper-V instances.

Troubleshooting layer-2 switches

Issues may occur because of the way an HA cluster assigns MAC addresses to the primary unit. Two clusters with the same group ID cannot connect to the same switch and cannot be installed on the same network unless they are separated by a router.

Forwarding delay on layer 2 switches

You must ensure that if there is a switch between the FortiGate HA cluster and the network it is protecting and the switch has a forwarding delay (even if spanning tree is disabled) when one of its interfaces is activated then the forwarding delay should be set as low as possible. For example, some versions of Cisco IOS have a forwarding delay of 15 seconds even when spanning tree is disabled. If left at this default value then TCP session pickup can fail because traffic is not forwarded through the switch on HA failover.

Failover issues with layer-3 switches

After a failover, the new primary unit sends gratuitous ARP packets to refresh the MAC forwarding tables of the switches connected to the cluster. If the cluster is connected using layer-2 switches, the MAC forwarding tables (also called arp tables) are refreshed by the gratuitous ARP packets and the switches start directing packets to the new primary unit.

In some configurations that use layer-3 switches, after a failover, the layer-3 switches may not successfully re-direct traffic to the new primary unit. The possible reason for this is that the layer-3 switch might keep a table of IP addresses and interfaces and may not update this table for a relatively long time after the failover (the table is not updated by the gratuitous ARP packets). Until the table is updated, the layer-3 switch keeps forwarding packets to the now failed cluster unit. As a result, traffic stops and the cluster does not function.

As of the release date of this document, Fortinet has not developed a workaround for this problem. One possible solution would be to clear the forwarding table on the layer-3 switch.

The `config system ha link-failed-signal` command described in [Updating MAC forwarding tables when a link failover occurs on page 1553](#) can be used to resolve link failover issues similar to those described here.

Failover and attached network equipment

It normally takes a cluster approximately 6 seconds to complete a failover. However, the actual failover time may depend on how quickly the switches connected to the cluster interfaces accept the cluster MAC address update from the primary unit. If the switches do not recognize and accept the gratuitous ARP packets and update their MAC forwarding table, the failover time will increase.

Also, individual session failover depends on whether the cluster is operating in active-active or active-passive mode, and whether the content of the traffic is to be virus scanned. Depending on application behavior, it may take a TCP session a longer period of time (up to 30 seconds) to recover completely.

Ethertype conflicts with third-party switches

Some third-party network equipment may use packets with Ethertypes that are the same as the ethertypes used for HA heartbeat packets. For example, Cisco N5K/Nexus switches use Ethertype 0x8890 for some functions. When one of these switches receives Ethertype 0x8890 heartbeat packets from an attached cluster unit, the switch generates CRC errors and the packets are not forwarded. As a result, FortiGates connected with these switches cannot form a cluster.

In some cases, if the heartbeat interfaces are connected and configured so regular traffic flows but heartbeat traffic is not forwarded, you can change the configuration of the switch that connects the HA heartbeat interfaces to allow level2 frames with Ethertypes 0x8890, 0x8891, and 0x8893 to pass.

You can also use the following CLI commands to change the Ethertypes of the HA heartbeat packets:

```
config system ha
  set ha-eth-type <ha_ethertype_4-digit_hex>
  set hc-eth-type <hc_ethertype_4-digit_hex>
  set l2ep-eth-type <l2ep_ethertype_4-digit_hex>
end
```

For more information, see [Heartbeat packet Ethertypes on page 1527](#).

LACP, 802.3ad aggregation and third-party switches

If a cluster contains 802.3ad aggregated interfaces you should connect the cluster to switches that support configuring multiple Link Aggregation (LAG) groups.

The primary and subordinate unit interfaces have the same MAC address, so if you cannot configure multiple LAG groups a switch may place all interfaces with the same MAC address into the same LAG group; disrupting the operation of the cluster.

You can change the FortiGate configuration to prevent subordinate units from participating in LACP negotiation. For example, use the following command to do this for an aggregate interface named Port1_Port2:

```
config system interface
  edit Port1_Port2
    set lacp-ha-slave disable
  end
```

This configuration prevents the subordinate unit interfaces from sending or receiving packets. Resulting in the cluster not being able to operate in active-active mode. As well, failover may be slower because after a failover the new primary unit has to perform LACP negotiation before being able to process network traffic.

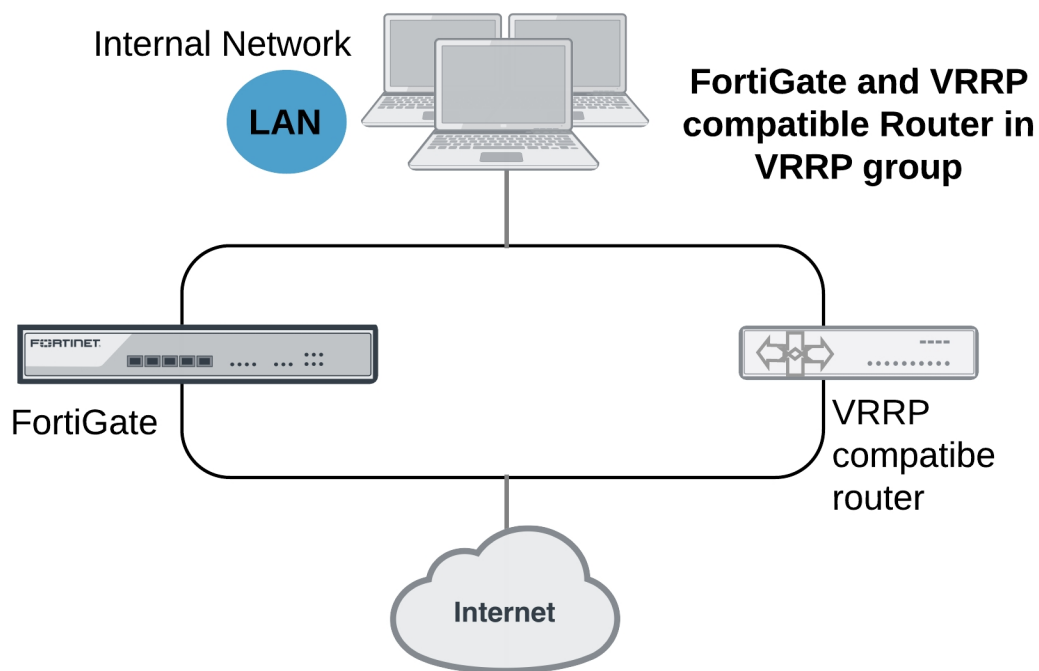
For more information, see [FGCP HA with 802.3ad aggregated interfaces on page 1436](#).

VRRP high availability

A Virtual Router Redundancy Protocol (VRRP) configuration can be used as a high availability solution to make sure that a network maintains connectivity with the internet (or with other networks) even if the default router for the network fails. Using VRRP, if a router or a FortiGate fails, all traffic to this router transparently fails over to another router or FortiGate that takes over the role of the router or FortiGate that failed. If the failed router or FortiGate is restored, it will once again take over processing traffic for the network. VRRP is described by [RFC 3768](#).

FortiOS supports VRRP versions 2 and 3 and you can set up VRRP groups that include multiple FortiGates and with other VRRP compatible routers. You can add different FortiGate models to the same VRRP group. FortiOS supports IPv4 and IPv6 VRRP and you can add IPv4 and IPv6 VRRP virtual routers to the same interface. FortiGates can also be quickly and easily integrated into a network that has already deployed a group of routers using VRRP.

Example VRRP configuration



The most common application of VRRP is to provide redundant default routers between an internal network and the internet. The default routers can be FortiGates and or any routers that support VRRP.

To set up VRRP:

1. Add a virtual VRRP router to the internal interface of each of the FortiGates and routers. This adds the FortiGates and routers to the same VRRP group.
2. Set the VRRP IP address of the group to the internal network default gateway IP address.

3. Give one of the VRRP group members the highest priority so it becomes the primary (or master) router and give the others lower priorities so they become backup routers.

During normal operations, all traffic from the internal network to the internet passes through the primary VRRP router. The primary router also sends VRRP advertisement messages to the backup routers. A backup router will not attempt to become a primary router while receiving these messages. If the primary router fails, the backup router with the highest priority becomes the new primary router after a short delay. During this delay the new primary router sends gratuitous ARP packets to the network to map the network's default route IP address to the new primary router's MAC address. All packets sent to the default route are now sent the new primary router. If the new primary router is a FortiGate, the network continues to benefit from FortiOS security features. If the new primary router is just a router, traffic continues to flow, but FortiOS security features are unavailable until the FortiGate is back on line.

If the backup router is a FortiGate, during a VRRP failover, as the FortiGate begins operating as the new primary router it will not have session information for all of the failed over in-progress sessions. So it would normally not be able to forward in-progress session traffic. To resolve this problem, immediately after a failover and for a short time (called the start time) the FortiGate acting as the new primary router operates with asymmetric routing enabled. This allows it to re-create all of the in-progress sessions and add them to its session table.

While operating with asymmetric routing enabled, the FortiGate cannot apply security functions. When the start-time ends the FortiGate disables asymmetric routing and returns to normal operation (including applying security functions).

Configuring VRRP

To configure VRRP you must configure two or more FortiGate interfaces or routers with the same virtual router ID and IP address. Then these FortiGates or routers can automatically join the same VRRP group. You must also assign priorities to each of the FortiGates or routers in the VRRP group. One of the FortiGates or routers must have the highest priority to become the primary (or master) router. The other FortiGates or routers in the group are assigned lower priorities and become backups. All of the routers in the VRRP group should have different priorities. If the primary router fails, VRRP automatically fails over to the router in the group with the next highest priority.

You configure VRRP from the FortiGate CLI by adding a VRRP virtual router to a FortiGate interface. You can add VRRP virtual routers to multiple FortiGate interfaces and you can add more than one virtual router to the same interface.

FortiOS supports VRRP with IPv6. You can add a VRRP virtual router to any FortiGate interface. You can also add IPv4 and IPv6 virtual routers to the same interface.

Adding an IPv4 VRRP virtual router to a FortiGate interface

Use the following command to add an IPv4 VRRP virtual router to the port10 interface of a FortiGate. This VRRP virtual router has a virtual router ID of 200, uses IP address 10.31.101.200 and has a priority of 255. Since this is the highest priority this interface is configured to be the master of the VRRP group with ID number 200.



VRRP can only be configured on physical interfaces or VLAN interfaces. You cannot configure VRRP on hardware-switch interfaces where multiple physical interfaces are combined into a hardware switch interface.

```
config system interface
edit port10
config vrrp
```

```
edit 200
    set vrip 10.31.101.200
    set priority 255
end
end
```

Adding an IPv6 VRRP virtual router to a FortiGate interface

Use the following command to add an IPv6 VRRP virtual router to the port20 interface of a FortiGate. This VRRP virtual router has a virtual router ID of 220, uses IP address 2001:db8:1::12 and has a priority of 255. Since this is the highest priority this interface is configured to be the master of the VRRP group with ID number 220.

```
config system interface
    edit port20
        config ipv6
            config vrrp6
                edit 220
                    set vrip 2001:db8:1::12
                    set priority 255
                end
            end
        end
    end
```

Setting up VRRP failover

VRRP routers in a VRRP group periodically send VRRP advertisement messages to all of the routers in the group to maintain one router as the primary router and the others as backup routers. The primary router is the one with the highest priority. If the backup routers stop receiving these packets from the primary router, the backup router with the highest priority becomes the new primary router.

The primary router stops sending VRRP advertisement messages if it either fails or becomes disconnected. You can also configure VRRP destination addresses that the primary router monitors. If the primary router becomes unable to connect to these destination addresses, it stops sending VRRP advertisement messages and the backup router with the highest priority becomes the primary router. You can add one or two destination addresses to a VRRP configuration. To be most effective, these destination addresses should be remote addresses.

For example, configure IPv4 VRRP on port14 with two destination address:

```
config system interface
    edit port14
        config vrrp
            edit 12
                set vrdst 10.10.10.20 10.20.20.10
            end
        end
    end
```

Configure IPv6 VRRP on port23 with one destination address:

```
config system interface
    edit port23
        config ipv6
            config vrrp6
                edit 223
                    set vrdst 2001:db8:1::12
                end
            end
        end
    end
```

IPv4 VRRP active failover

You can reduce IPv4 VRRP failover times with the `vrdst-priority` option. This option causes the primary router to actively signal to the backup routers when the primary router can't reach its configured destination address or addresses. The primary router does this by sending a lower priority for itself in the VRRP advertisement messages. You set this lower priority with the `vrdst-priority` option. The backup router with the highest priority becomes the new primary router and takes over processing traffic.

The following example configures the primary router to have a priority of 255 so it should always become the primary router. The command also sets `vrdst-priority` to 10. So if the primary router can no longer connect to its destination address of 10.10.10.1, the primary router informs the VRRP group that its priority is now 10.

```
config system interface
  edit port10
    config vrrp
      edit 12
        set vrip 10.31.101.200
        set priority 255
        set vrdst 10.10.10.1
        set vrdst-priority 10
      end
    end
```

Failover of IPv4 firewall VIPs and IP pools

FortiOS VRRP HA supports failover of firewall VIPs and IP Pools when the status of a virtual router (VR) changes. This feature introduces a new proxy ARP setting to map VIP and IP Pool address ranges to each VR's Virtual MAC (VMAC). After failover, the IP Ranges added to the new primary VR are routed to the new primary VR's VMAC.

Use the following command to add a proxy ARP address range and a single IP address to a VR added to a FortiGate's port5 interface. The address range and single IP address should match the address range or single IP for VIPs or IP Pools added to the port5 interface:

```
config system interface
  edit port5
    config vrrp
      edit 1
        config proxy-arp
          edit 1
            set ip 192.168.62.100-192.168.62.200
          next
          edit 2
            set ip 192.168.62.225
          end
        end
      end
    end
```

Changing the advertisement message interval

By default, VRRP advertisement messages are sent once a second. You can use the following to change the frequency of sending these messages. The range is 1 to 255 seconds.

For example, configure an IPv4 VRRP to send advertisement messages every 10 seconds:

```
config system interface
  edit port14
    config vrrp
      edit 12
        set adv-interval 10
      end
    end
```

Configure IPv6 VRRP to send advertisement messages every 20 seconds:

```
config system interface
  edit port23
    config ipv6
      config vrrp6
        edit 223
          set adv-interval 20
        end
      end
    end
```

Changing how long backup routers wait before assuming the primary router has failed

The VRRP startup time is the maximum time that a backup router waits between receiving advertisement messages from the primary router. If the backup router has to wait longer than the start-time, it assumes the primary router has failed and becomes the new primary router.

The default startup time is 3 seconds and the range is 1 to 255 seconds.

In some cases the advertisement messages may be delayed. For example, some switches with spanning tree enabled may delay some of the advertisement message packets. If you find that backup routers are attempting to become primary routers even though the primary router hasn't failed, you can extend the start time to make sure the backup routers wait long enough for the advertisement messages.

For example, set the IPv4 VRRP startup time to 10 seconds:

```
config system interface
  edit port14
    config vrrp
      edit 12
        set start-time 10
      end
    end
```

Configure set the IPv6 VRRP startup time to 15 seconds:

```
config system interface
  edit port23
    config ipv6
      config vrrp6
        edit 223
          set start-time 15
        end
      end
    end
```

Using VRRP virtual MAC addresses

The VRRP virtual MAC address (or virtual router MAC address) is a shared MAC address adopted by the primary router. If the primary router fails, the same virtual MAC address is picked up by the new primary router allowing all devices on the network to transparently connect to the default route using the same virtual MAC address. You must enable the VRRP virtual MAC address feature on all members of a VRRP group.

Each VRRP router is associated with its own virtual MAC address. The last part of the virtual MAC depends on the VRRP virtual router ID using the following format:

```
00-00-5E-00-01-<VRID_hex>
```

Where <VRID_hex> is the VRRP virtual router ID in hexadecimal format in internet standard bit-order. For more information about the format of the virtual MAC see [RFC 3768](#).

Some examples:

- If the VRRP virtual router ID is 10 the virtual MAC would be 00-00-5E-00-01-0a.
- If the VRRP virtual router ID is 200 the virtual MAC would be 00-00-5E-00-01-c8.

The VRRP virtual MAC address feature is disabled by default. When you enable the feature on a FortiGate interface, all of the VRRP routers added to that interface use their own VRRP virtual MAC address. Each virtual MAC address will be different because each virtual router has its own ID.

Use the following command to enable the VRRP virtual MAC address for an IPv4 VRRP configuration on the port2 interface:

```
config system interface
  edit port2
    set vrrp-virtual-mac enable
  end
end
```

The port2 interface will now accept packets sent to the MAC addresses of the IPv4 VRRP virtual routers added to this interface.

Use the following command to enable the VRRP virtual MAC address for an IPv6 VRRP configuration on the port22 interface:

```
config system interface
  edit port22
    config ipv6
      set vrrp-virtual-mac6 enable
    end
  end
end
```

The port22 interface now accepts packets sent to the MAC addresses of the IPv6 VRRP virtual routers added to this interface.

Since devices on the LAN do not have to learn a new MAC address for a new VRRP router in the event of a failover, this feature can improve network efficiency, especially on large and complex networks.

If the VRRP virtual MAC address feature is disabled, the VRRP group uses the MAC address of the master. In the case of a FortiGate VRRP virtual router this is the MAC address of the FortiGate interface that the VRRP virtual routers are added to. If a master fails, when the new master takes over it sends gratuitous ARPs to associate the VRRP virtual router IP address with the MAC address of the new master (or the interface of the FortiGate that has become the new master). If the VRRP virtual MAC address is enabled the new master uses the same MAC address as the old master.

Setting up VRRP groups

A VRRP group includes all the relevant VRRP IDs and tracks the VRRP status to force the status of all group members if a VRRP domain is changed from primary to backup. The VRRP group ID can be between 1 and 65535.

For an IPv4 VRRP configuration, use the following command to add a VRRP group to the port20 interface that includes the virtual route identifiers 25, 50, 66 and 70 to VRRP group 10.

```
config system interface
  edit port20
    config vrrp
      edit 25
```

```
        set vrgrp 10
    next
    edit 50
        set vrgrp 10
    next
    edit 66
        set vrgrp 10
    next
    edit 70
        set vrgrp 10
    end
```

For an IPv6 VRRP configuration, use the following command to add a VRRP group to the port22 interface that includes the virtual route identifiers 10, 51, and 83 to VRRP group 60.

```
config system interface
    edit port22
        config ipv6
            config vrrp6
                edit 10
                    set vrgrp 60
                next
                edit 51
                    set vrgrp 60
                next
                edit 83
                    set vrgrp 60
            end
```


Example IPv4 VRRP configuration: two FortiGates in a VRRP group

This example includes a VRRP group consisting of two FortiGates that connect an internal network to the internet. As shown below, the internal network's default route is 10.31.101.120.

The FortiGate port2 interfaces connect to the internal network. A VRRP virtual router is added to each FortiGate's port2 interface. The virtual router IP address is 10.31.101.120 (the internal network's default route) and the virtual router's ID is 5. The VRRP priority of the primary router is set to 255 and the VRRP priority of the backup router is 50. The port2 interface of each FortiGate should have an IP address that is different from the virtual router IP address and the port2 interface IP addresses should be different from each other.

This example also includes enabling the VRRP virtual MAC address on both FortiGate port2 interfaces so that the VRRP group uses the VRRP virtual MAC address.

Example VRRP configuration with two FortiGates



To configure the FortiGates for VRRP

1. Select one of the FortiGates to be the primary VRRP router and the other to be the backup router.
2. From the primary router CLI, enter the following command to enable the VRRP virtual MAC address on the port2 interface and add the VRRP virtual router to the port2 interface:

```
config system interface
  edit port2
    set vrrp-virtual-mac enable
    config vrrp
      edit 5
        set vrip 10.31.101.120
        set priority 255
      end
    end
  end
```

3. From the backup router CLI, enter the following command to enable the VRRP virtual MAC address on the port2 interface and add the VRRP virtual router to the port2 interface:

```
config system interface
  edit port2
    set vrrp-virtual-mac enable
    config vrrp
      edit 5
        set vrip 10.31.101.120
        set priority 50
      end
    end
  end
```

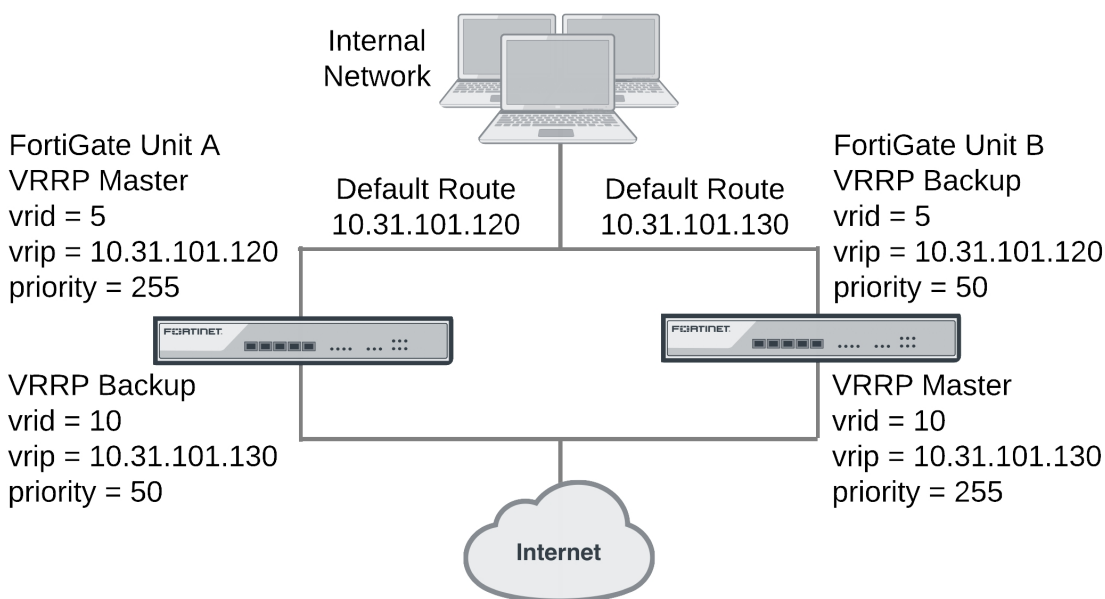
Example IPv4 VRRP configuration: VRRP load balancing two FortiGates and two VRRP groups

In this configuration two VRRP groups are involved. Each FortiGate participates in both of them. One FortiGate is the primary router of one group and the other FortiGate is the primary router of the other group. The network distributes traffic between two different default routes (10.31.101.120 and 10.31.101.130). One VRRP group is configured with one of the default route IP addresses and the other VRRP group gets the other default route IP address. During normal operation, both FortiGates are processing traffic and the VRRP groups are used to load balance the traffic between the two FortiGates.

If one of the FortiGates fails, the remaining FortiGate becomes the primary router of both VRRP groups. The network sends all traffic for both default routes to this FortiGate. The result is a configuration that, under normal operation load, balances traffic between two FortiGates, but if one of the FortiGates fails, all traffic fails over to the FortiGate that is still operating.

This example also includes enabling the VRRP virtual MAC address on both FortiGate port2 interfaces so that the VRRP groups use their VRRP virtual MAC addresses.

Example VRRP configuration with two FortiGates and two VRRP groups



To configure the FortiGates

1. Log into the CLI of FortiGate A.
2. Enter the following to enable the VRRP virtual MAC address feature and add the VRRP groups to the port2 interface of FortiGate A:

```
config system interface
edit port2
set vrrp-virtual-mac enable
config vrrp
```

```
edit 50 (32)
    set vrip 10.31.101.120
    set priority 255
next
edit 100 (64)
    set vrip 10.31.101.130
    set priority 50
end
end
```

3. Log into the CLI of FortiGate B.

4. Enter the following command to enable the VRRP virtual MAC address feature and add the VRRP groups to the port2 interface of FortiGate B:

```
config system interface
edit port2
    set vrrp-virtual-mac enable
config vrrp
edit 50
    set vrip 10.31.101.120
    set priority 50
next
edit 100
    set vrip 10.31.101.130
    set priority 255
end
end
```

Optional VRRP configuration settings

In addition to the basic configuration settings, you can change to the VRRP configuration in the following ways. All of these options apply to both IPv4 and IPv6 VRRP unless noted.

- Enable or disable individual virtual router configurations using the `status` option. Normally virtual router configurations are enabled but you can temporarily disable one if it is not required.
- Enable or disable preempt mode using the `preempt` option. In preempt mode, a higher priority backup router can preempt a lower priority primary router. This can happen if the primary router has failed, a backup router has become the primary router, and the failed primary router restarts. Since the restarted router has a higher priority, if preempt mode is enabled the restarted router replaces the current primary router becoming the new primary router. Preempt mode is enabled by default.
- You can add one or two destination addresses (`vrdst`) to a VRRP configuration. To be most effective, these destination addresses should be remote addresses.

FortiController-5000 SLBC support

FortiController-5000 Session Aware Load Balancing Clustering (SLBC) supports the Virtual Router Redundancy Protocol (VRRP), allowing you to configure HA between FortiController-5000 SLBC clusters using VRRP. You can also add a FortiController-5000 SLBC cluster to a VRRP group with other VRRP routers.

You configure VRRP on the FortiController-5000 by creating a VRRP group and adding one or more FortiController front panel interfaces to the group. During normal operation, the primary FortiController sends outgoing VRRP routing advertisements. Both the primary and backup FortiControllers listen for incoming VRRP advertisements from other routers in the VRRP group. If the primary FortiController fails, the new primary

FortiController takes over the role of both sending and receiving VRRP advertisements, maintaining the FortiController-5000 cluster within the VRRP group.

FortiGate Session Life Support Protocol (FGSP)

In a network that already includes load balancing (either with load balancers or routers) for traffic redundancy, two or more FortiGates can be integrated into the load balancing configuration using the FortiGate Session Life Support Protocol (FGSP). The external load balancers or routers can distribute sessions among the FortiGates and the FGSP performs session synchronization of IPv4 and IPv6 TCP, SCTP, UDP, ICMP, expectation, and NAT sessions and IPsec tunnels to keep the session tables of the FortiGates synchronized. If one of the FortiGates fails, session failover occurs and active sessions fail over to the FortiGates that are still operating. This failover occurs without any loss of data. As well, the external routers or load balancers will detect the failover and re-distribute all sessions to the peers that are still operating.

The FortiGates operate as peers that process traffic and synchronize sessions with other FortiGates in the cluster. An FGSP cluster can include from 2 to 16 FortiGates. Adding more FortiGates increases the CPU and memory required to keep all of the FortiGates synchronized. So depending on your network conditions, adding too many FortiGates to an FGSP cluster may reduce overall performance.

The FortiGates in the FGSP cluster must be the same model and be running the same firmware version. You use the `config system cluster-sync` command to configure FGSP between the FortiGates and the `config system ha` command to configure what is synchronized.

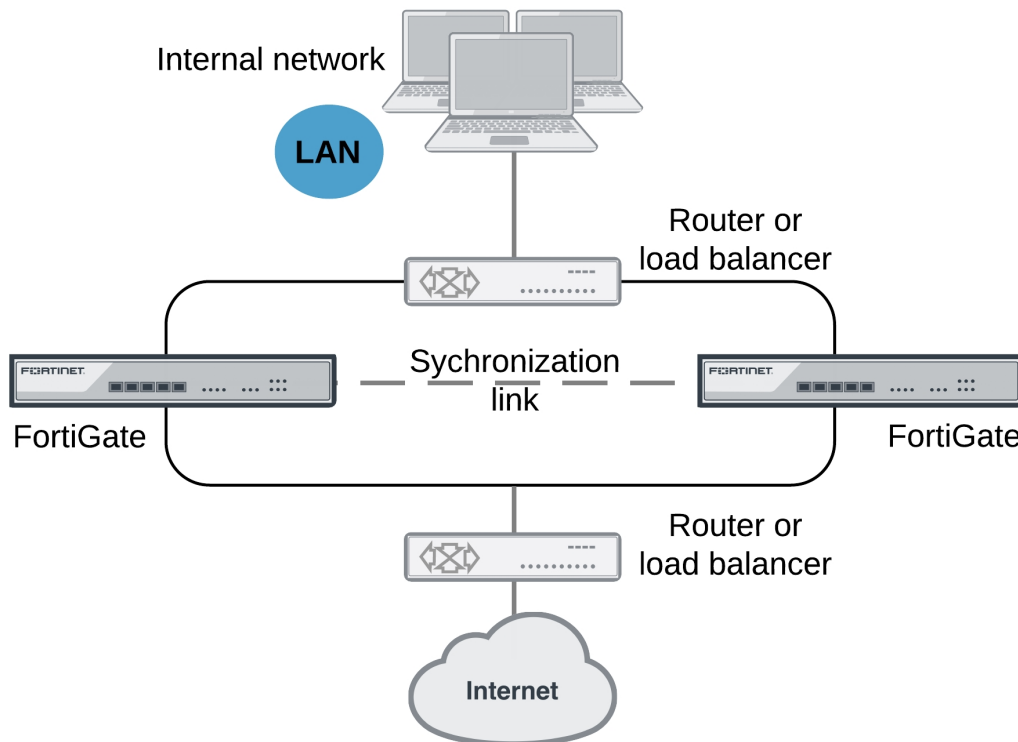


In previous versions of FortiOS the FGSP was called TCP session synchronization or standalone session synchronization. The FGSP has been expanded to include configuration synchronization and session synchronization of connectionless sessions, expectation sessions, and NAT sessions and IPsec tunnels.



FGSP also supports synchronizing sessions between FGCP clusters and between FGCP clusters and standalone FortiGates. See [Session synchronization between FGCP clusters on page 1608](#). FGSP is also compatible with FortiGate VRRP.

The FGSP can be used instead of FGCP HA to provide **session synchronization** between two peer FortiGates. If the external load balancers direct all sessions to one peer the affect is similar to active-passive FGCP HA. If external load balancers or routers load balance traffic to both peers, the effect is similar to active-active FGCP HA. The load balancers should be configured so that all of the packets for any given session are processed by the same peer. This includes return packets.

FGSP HA

By default, FGSP synchronizes all IPv4 and IPv6 TCP sessions, IPsec tunnels, and also synchronizes the configuration of the FortiGates.

You can optionally enable session pickup to synchronize connectionless (UDP and ICMP) sessions, expectation sessions, and NAT sessions. If you do not enable session pickup, the FGSP does not share session tables for the particular session type and sessions do not resume after a failover. All sessions are interrupted by the failover and must be re-established at the application level. Many protocols can successfully restart sessions with little, or no, loss of data. Others may not recover easily. Enable session pickup for sessions that may be difficult to reestablish. Since session pickup requires FortiGate memory and CPU resources, only enable this feature for sessions that you need to have synchronized.

The synchronization link is set up in the same way as FGCP heartbeat interfaces. You must connect the synchronization link interfaces together and use the heartbeat device (hbdev) option to add the heartbeat devices to the configuration.

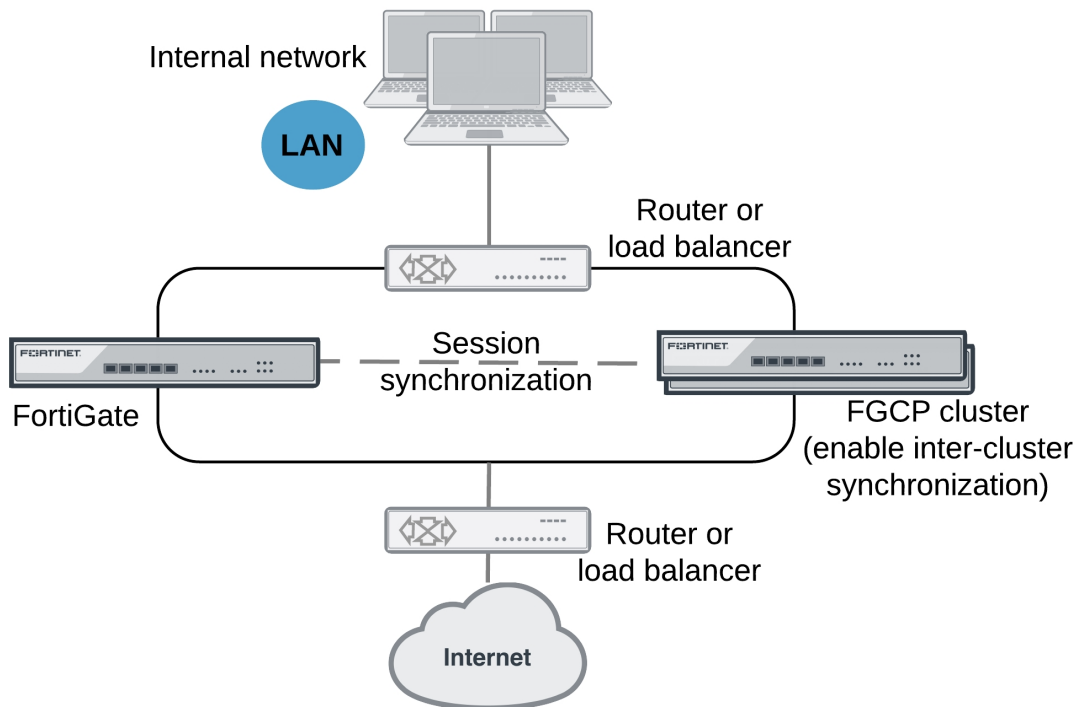
You can also optionally add filters to control which sessions are synchronized. You can add filters to only synchronize packets from specified source and destination addresses, specified source and destination interfaces, and specified services.

Load balancing and session failover is done by external routers or load balancers instead of by the FGSP. The FortiGates only perform session synchronization to support session failover as well as configuration synchronization.

Session synchronization between FGCP clusters

Session synchronization between FGCP clusters (also called inter-cluster session synchronization) allows you to synchronize sessions among FGCP clusters and standalone FortiGates. The FGSP can synchronize sessions among up to four FGCP clusters and FortiGates.

Example session synchronization between a FortiGate and an FGCP cluster



Enter the following command to enable inter-cluster synchronization on an FGCP cluster:

```
config system ha
  set inter-cluster-session-sync enable
end
```

Once you enable inter-cluster session synchronization, all FGSP configuration options are available in the FGCP cluster CLI and you can set up the FGSP configuration in the same way as for standalone FortiGates.

Inter-cluster session synchronization is compatible with all FGCP operating modes, such as active-active, active-passive, virtual clustering, full mesh HA.

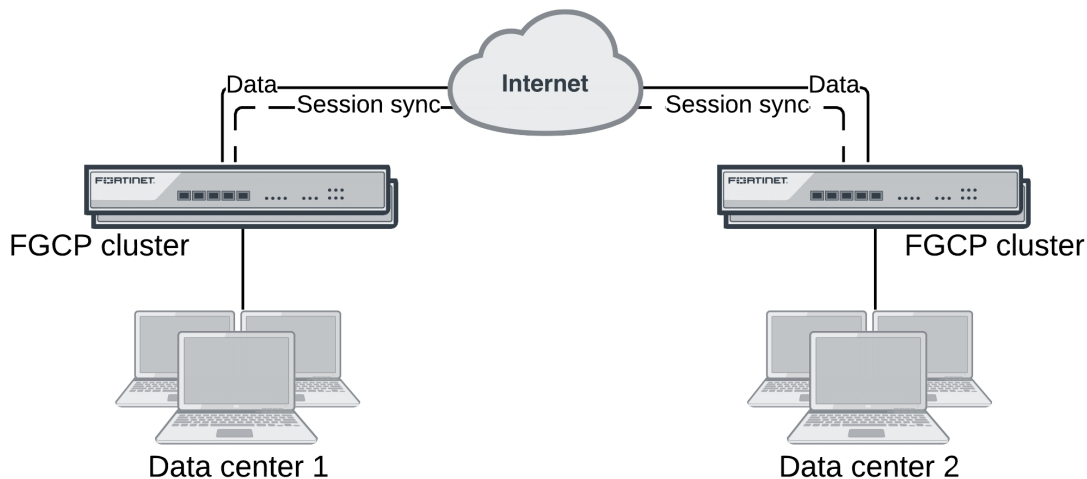
What is synchronized?

Inter-cluster session synchronization synchronizes all supported FGSP session types, including TCP sessions, IPsec tunnels, IKE routes, connectionless sessions (UDP and ICMP), NAT sessions, asymmetric sessions, and expectation sessions. Inter-cluster session synchronization doesn't support configuration synchronization.

Inter-cluster synchronization between data centers

Inter-cluster session synchronization is deployed for session-synchronization among multiple data centers if one or more of the data centers is protected by an FGCP cluster.

Example inter-cluster session synchronization between two data centers



In this example, you enable inter-cluster session synchronization for both of the clusters, and then configure session synchronization options on each cluster, as required.

Configuring FGSP HA cluster-sync instances

You use the following command to configure an FGSP HA cluster-sync instance.

```
config system cluster-sync
  edit 1
    set peerip <peer-ip-address>
    set peervd <vdom-name>
    set syncvd <vdom-name>
  end
```

Where:

- `peerip` is the IP address of an interface of another FortiGate in the FGSP cluster that this configuration synchronizes sessions to.
- `peervd` is the name of the VDOM on the other FortiGate that should be synchronized with this one. By default the `peervd` is `root`.
- `syncvd` is the name of the VDOM of the FortiGate that should be synchronized with the other FortiGate. If multiple VDOMs are not enabled, `syncvd` should be set to `root`.



For FGSP HA to work properly, all VDOMs to be synchronized must be added to all of the FortiGates in cluster. The names of the matching interfaces in each VDOM must also be the same; this includes the names of matching VLAN interfaces. Note that the index numbers of the matching interfaces and VLAN interfaces can be different. Also the VLAN IDs of the matching VLAN interfaces can be different. If you enable configuration synchronization this will happen automatically.

This command creates a cluster-sync instance that causes a FortiGate to synchronize the TCP sessions of one of its VDOMs (by default the root VDOM) to the root VDOM of another FortiGate (which would become another FortiGate in the FGSP cluster). You can also use the `config system ha` command to synchronize more session types and to synchronize the configuration. Cluster-sync instances are not synchronized and must be added to each FortiGate in the cluster.

A cluster of two FortiGates would only require one `cluster-sync` instance for each VDOM to be synchronized. This instance would synchronize the sessions from the root VDOM of one FortiGate to the root VDOM of the other. The second FortiGate would also include a cluster-sync instance to synchronize its root VDOM with the root VDOM of the other FortiGate.

In a multiple VDOM configuration, you add a separate cluster-sync instance for each VDOM to be synchronized. You don't have to synchronize all VDOMs. If multiple VDOMs are enabled, the `config system cluster-sync` command is a global command.

FGSP clusters with three or more FortiGates

If an FGSP cluster includes three or more FortiGates you must explicitly define all of the cluster-sync instances that you need. In a cluster of four FortiGates, each FortiGate can be synchronized with up to three other FortiGates. So, to synchronize all of the FortiGates, you must add three cluster-sync instances to each FortiGate (or $n-1$, where n is the number of FortiGates in the cluster).

Selecting the sessions to synchronize

You can add a cluster-sync instance with a filter to only synchronize some sessions. A filter can be added to a cluster-sync instance as follows:

```
config system cluster-sync
  edit 1
    set peerip <peer-ip-address>
    set peervd <vdom-name>
    set syncvd <vdom-name>
    config session-sync-filter
      srcintf <interface-name>
      dstintf <interface-name>
      srcaddr x.x.x.x x.x.x.x
      dstaddr x.x.x.x x.x.x.x
      srcaddr6 ::/x
      dstaddr6 ::/x
    end
  end
```

You can use the filter to only synchronize sessions according to the session source and destination interface and IPv4 or IPv6 address.

You can only add one filter to a cluster-sync instance. To create multiple filters you must create multiple cluster-sync instances.

Synchronizing TCP and SCTP sessions

Use the following to enable session synchronization for TCP and SCTP sessions and to configure the FGSP to use the port8 interface for synchronizing traffic:

```
config system ha
    set session-pickup enable
    set hbdev "port8" 50
end
```

Automatic session sync after peer reboot

You can configure your FGSP cluster to resume sessions more smoothly after a failed FortiGate rejoins the cluster. In some cases when a failed FortiGate in the cluster comes back up it may begin processing sessions before the session table from the other FortiGate in the cluster has been synchronized to it. When this happens, the FortiGate may drop packets until the session synchronization is complete.

Shutting down interfaces during session synchronization

This feature allows you to shut down some interfaces on the failed FortiGate when it is starting up so that it will not accept packets until session synchronization is complete. Then the interfaces are brought up and traffic can flow. While the interfaces are down, the FortiGate that had not failed keeps processing traffic.

Use the following to select the interfaces to shutdown while waiting for session synchronization to complete:

```
config system cluster-sync
    edit 1
        set down-intfs-before-sess-sync port1 port2
    end
```

Heartbeat monitoring

If the FortiGate that was running fails before session synchronization is complete, the FortiGate that is restarting will not be able to complete session synchronization and will not turn on its shutdown interfaces. To prevent this from happening, FGSP includes heartbeat monitoring. Using heartbeat monitoring, the FortiGate that is waiting for session synchronization to finish can detect that the other FortiGate is down and turn on its interfaces even if session synchronization is not complete. You can use the following to change the heartbeat interval (`hb-interval`) and lost heartbeat threshold (`hb-lost-threshold`) to change heartbeat monitoring timing.

```
config system cluster-sync
    edit 1
        set hb-interval 2
        set hb-lost-threshold 3
    end
```

Synchronizing the configuration

The FGSP includes configuration synchronization, allowing you to make configuration changes once for all of the FortiGates in the cluster, instead of requiring you to make redundant configuration changes on each FortiGate.

By default, configuration synchronization is disabled. You can enter the following to enable it:

```
config system ha
    set standalone-config-sync enable
end
```

You must enter this command on all of the FortiGates in the FGSP cluster. When you enable synchronizing the configuration, the FGSP uses FGCP primary unit selection to select a config sync primary (or master) FortiGate. The other FortiGates in the FGSP cluster become config sync backup FortiGates. The FGSP synchronizes all configuration changes that you make on the config sync primary FortiGate to the config sync backup FortiGates. Fortinet recommends making all configuration changes on the config sync primary FortiGate.

Config sync primary FortiGate selection

The FGSP cluster uses FGCP primary unit selection to select the FGSP config sync primary FortiGate (see [Primary unit selection on page 1385](#)). So normally, the FortiGate with the highest serial number would become the primary FortiGate.

You can use device priority select one of the FortiGates to become the config sync primary FortiGate. For example, the following command enables configuration synchronization on a FortiGate and sets a higher device priority than the default of 128 to make sure that this FortiGate becomes the primary FortiGate.

```
config system ha
    set standalone-config-sync enable
    set priority 250
end
```

Settings that are not synchronized

FGSP configuration synchronization does not synchronize settings that identify the FortiGate to the network. The following settings are not synchronized:

- Transparent mode management IPv4 and IPv6 IP addresses and default gateways.
- All `config system cluster-sync` settings.
- All `config system interface` settings except `vdom`, `vlanid`, `type` and `interface`.
- All `config firewall sniffer` settings.
- All router BFD and BFD6 settings.
- The following BGP settings: `as`, `router-id`, `aggregate-address`, `aggregate-address6`, `neighbor-group`, `neighbor`, `network`, and `network6`.
- The following OSPF settings: `router-id`, `area`, `ospf-interface`, `network`, `neighbor`, and `summary-address`.
- The following OSPF6 settings: `router-id`, `area`, and `ospf6-interface`.
- All RIP settings.
- All policy routing settings.
- All static routing settings.

FGSP and firmware upgrades

The steps to follow to upgrade the firmware running on an FGSP cluster depend on whether you have enabled configuration synchronization (see [Synchronizing the configuration on page 1611](#)) or not:

- If you have not enabled configuration synchronization, you must upgrade the firmware separately on each FortiGate in the cluster. Upgrading the firmware of each FortiGate interrupts traffic through that FortiGate.
- If you have enabled configuration synchronization, you can upgrade the cluster firmware by upgrading the firmware running on the sync primary FortiGate. The FGSP then sends the new firmware image to the config sync backup FortiGates and all of the FortiGates in cluster install the new firmware and restart. The firmware upgrade simultaneously interrupts traffic through all of the FortiGates in the cluster.

Backing up and restoring the configuration of an FGSP cluster

You should maintain separate backup configuration files for each FortiGate in the FGSP cluster. When you restore the configuration of one of the FortiGates in an FGSP cluster, the FGSP does not synchronize the configuration file to the other FortiGates in the FGSP cluster. Instead you must back up each FortiGate separately and restore the configuration of each FortiGate separately with its own configuration file.

IPsec tunnel synchronization

When you use the `config system cluster-sync` command to enable FGSP, IPsec keys and other runtime data (but not actual tunnel sessions) are synchronized between cluster units. This means that if one of the cluster units goes down, the cluster units that are still operating can quickly get IPsec tunnels re-established without re-negotiating them. However, after a failover, all existing tunnel sessions on the failed FortiGate have to be restarted on the FortiGates that are still operating.

IPsec tunnel sync supports both static and dialup IPsec tunnels. For IPsec tunnel synchronization to work, the interfaces on the FortiGates that are tunnel endpoints must have the same IP addresses and external routers must be configured to load balance IPsec tunnel sessions to the FortiGates in the cluster.

Optionally synchronizing IKE routes

You can use the following command to control whether IKE routes are synchronized to all units in the FGSP cluster.

```
config system cluster-sync
  edit 0
    set slave-add-ike-routes {enable | disable}
  end
```

Enable to synchronize IKE routes, disable if you do not need to synchronize IKE routes. Enabling routing synchronization is optional but doing so increases synchronization overhead and bandwidth usage. If you have problems with IPsec VPN tunnel synchronization, you may want to enable synchronizing routes. Otherwise you could leave it disabled to improve performance and save bandwidth.

Synchronizing UDP and ICMP (connectionless) sessions

In many configurations, due to their non-stateful nature, UDP and ICMP sessions don't need to be synchronized to naturally failover. However, if required you can configure the FGSP to synchronize UDP and ICMP sessions by entering the following:

```
config system ha
  set session-pickup enable
  set session-pickup-connectionless enable
end
```

Synchronizing NAT sessions

By default, NAT sessions are not synchronized. However, the FGSP can synchronize NAT sessions if you enter the following:

```
config system ha
  set session-pickup enable
```

```
set session-pickup-nat enable
end
```

However, if you want NAT sessions to resume after a failover, you should not configure NAT to use the destination interface IP address since the FGSP FortiGates have different IP addresses. With this configuration, after a failover all sessions that include the IP addresses of interfaces on the failed FortiGate will have nowhere to go since the IP addresses of the failed FortiGate will no longer be on the network.

Instead, in an FGSP configuration, if you want NAT sessions to failover, you should use IP pools with the type set to overload (which is the default IP pool type). For example:

```
config firewall ippool
edit FGSP-pool
set type overload
set startip 172.20.120.10
set endip 172.20.120.20
end
```

Then when you configure NAT firewall policies, turn on NAT and select to use dynamic IP pool and select the IP pool that you added. Configuration synchronization should add the same IP pools and firewall policies to all FortiGates in the cluster. If configuration synchronization is not enabled you must add the same IP pools and policies to all of the FortiGates in the cluster.

Synchronizing asymmetric sessions

By default, asymmetric sessions are not synchronized. Normally, session synchronization cannot be asymmetric because it is stateful. So all of the packets of a given session must be processed on the same FortiGate. This includes return packets.

However, if you have an asymmetric routing configuration, you can enter the following command to synchronize asymmetric sessions by dynamically detecting asymmetric sessions and disabling anti-reply for these sessions.

```
config system ha
set session-pickup enable
set session-pickup-expectation enable
end
```

The FGSP enforces firewall policies for asymmetric traffic, including cases where the TCP 3-way handshake is split between two FortiGates. For example, FGT-A receives the TCP-SYN, FGT-B receives the TCP-SYN-ACK, and FGT-A receives the TCP-ACK. Under normal conditions, a firewall will drop this connection since the 3-way handshake was not seen by the same firewall. However, two FortiGates with FGSP configured will be able to properly pass this traffic since the firewall sessions are synchronized.

This asymmetric function can also work with connectionless UDP and ICMP traffic. If traffic will be highly asymmetric, as described above, the following command must be enabled on both FortiGates.

```
config system ha
set session-pickup enable
set session-pickup-connectionless enable
end
```

Synchronizing asymmetric traffic can be very useful in situations where multiple internet connections from different ISPs are spread across multiple FortiGates. Since it is typically not possible to guarantee internet-bound traffic leaving via an ISP will return using the exact same ISP, the FGSP provides critical firewall functions in this situation.



Asymmetric sessions may not be synchronized in low latency networks if the reply packet is received before the peer has received the session synchronization packet. This limitation usually only occurs in low latency networks.

The FGSP also has applications in virtualized computing environments where virtualized hosts move between data centers. The firewall session synchronization features of FGSP allow for more flexibility than in traditional firewalling functions.

Synchronizing expectation sessions

FortiOS session helpers keep track of the communication of Layer-7 protocols such as FTP and SIP that have control sessions and expectation sessions. Usually the control sessions establish the link between server and client and negotiate the ports and protocols that will be used for data communications. The session helpers then create expectation sessions through the FortiGate for the ports and protocols negotiated by the control session.

The expectation sessions are usually the sessions that actually communicate data. For FTP, the expectation sessions transmit files being uploaded or downloaded. For SIP, the expectation sessions transmit voice and video data. Expectation sessions usually have a timeout value of 30 seconds. If the communication from the server is not initiated within 30 seconds, the expectation session times out and traffic will be denied.

By default the FGSP does not synchronize expectation sessions and if a failover occurs, the sessions will have to be restarted.

If you want to synchronize expectation sessions so that they will continue after a failover, you can enter the following:

```
config system ha
    set session-pickup enable
    set session-pickup-expectation enable
end
```

GTP session synchronization: FGSP for FortiOS Carrier

FortiOS Carrier FGSP supports GTP session synchronization, allowing you to synchronize GTP sessions among up to four FortiGates or clusters running FortiOS Carrier. No special configuration is required to support GTP session synchronization. GTP sessions are synchronized like any other TCP or UDP session.

One limitation to this configuration is that asymmetric routing is not supported. So the router or load balancer that load balances the GTP sessions must send all of the packets of a session to the same FortiGate.

Also, for GTP sessions to be synchronized, session synchronization must be enabled. Also, since most GTP traffic uses UDP you should also enable UDP session synchronization. For example:

```
config system ha
    set session-pickup enable
    set session-pickup-connectionless enable
end
```

Security profile flow-based inspection and asymmetric traffic

Security profile inspection (flow- or proxy-based) for a session is not expected to work properly if the traffic in the session is balanced across more than one FortiGate in either direction. Flow-based inspection should be used in FGSP deployments.

For an environment where traffic is symmetric, security profile inspection can be used with the following limitations:

- No session synchronization for the sessions inspected using proxy-based inspection. Sessions will drop and need to be reestablished after data path failover.

- Sessions with flow-based inspection will failover; however, inspection of failed over sessions after the failover may not work.

A single FortiGate must see both the request and reply traffic for security profile inspection to function correctly. For environments where asymmetric traffic is expected, security profile inspection should not be used.

Notes and limitations

FGSP HA has the following limitations:

- The FGSP is a global configuration option. As a result you can only add one service to a filter configuration. You cannot add custom services or service groups even if virtual domains are not enabled.
- You can only add one filter configuration to a given FGSP configuration. However, you can add multiple filters by adding multiple identical FGSP configurations, each one with a different filter configuration.
- Sessions accepted by security policies with security profiles configured are not synchronized.
- FGSP HA is configured from the CLI.
- FGSP HA is available for FortiGates or virtual domains operating in NAT or transparent mode. NAT sessions are not synchronized in either mode (unless NAT synchronization is enabled as described in [Synchronizing NAT sessions on page 1613](#)). In NAT/Route mode, only sessions for route mode security policies are synchronized. In transparent mode, only sessions for normal transparent mode policies are synchronized.
- FGSP HA is supported for traffic on physical interfaces, VLAN interfaces, zones, aggregate interfaces, and NPx (NP4, NP6 etc.) accelerated interfaces. The FGSP has not been tested for inter-ldom links, between HA clusters, and for redundant interfaces.
- The names of the matching interfaces, including VLAN interfaces, aggregate interfaces and so on, must be the same on both peers.
- An FGSP cluster can include from 2 to 16 FortiGates. Adding more FortiGates increases the CPU and memory required to keep all of the FortiGates synchronized.

Configuring session synchronization links

When FGSP HA is operating, the FortiGates share session information over Ethernet links similar to an HA heartbeat link. Usually you would use the same interface on each FortiGate for session synchronization. If possible you should connect the session synchronization interfaces directly without using a switch or other networking equipment. For FortiGate-5000 systems you can use a backplane interface as the session synchronization link.

You can use different interfaces on each FortiGate for session synchronization links. Also, if you have multiple session synchronization configurations, you can have multiple links between the FortiGates. In fact if you are synchronizing a lot of sessions, you may want to configure and connect multiple session synchronization links to distribute session synchronization traffic to these multiple links.

You cannot configure backup session synchronization links. Each configuration only includes one session synchronization link.

The session synchronization link should always be maintained. If session synchronization communication is interrupted and a failure occurs, sessions will not failover and data could be lost.

Session synchronization traffic can use a considerable amount of network bandwidth. If possible, session synchronization link interfaces should only be used for session synchronization traffic and not for data traffic.

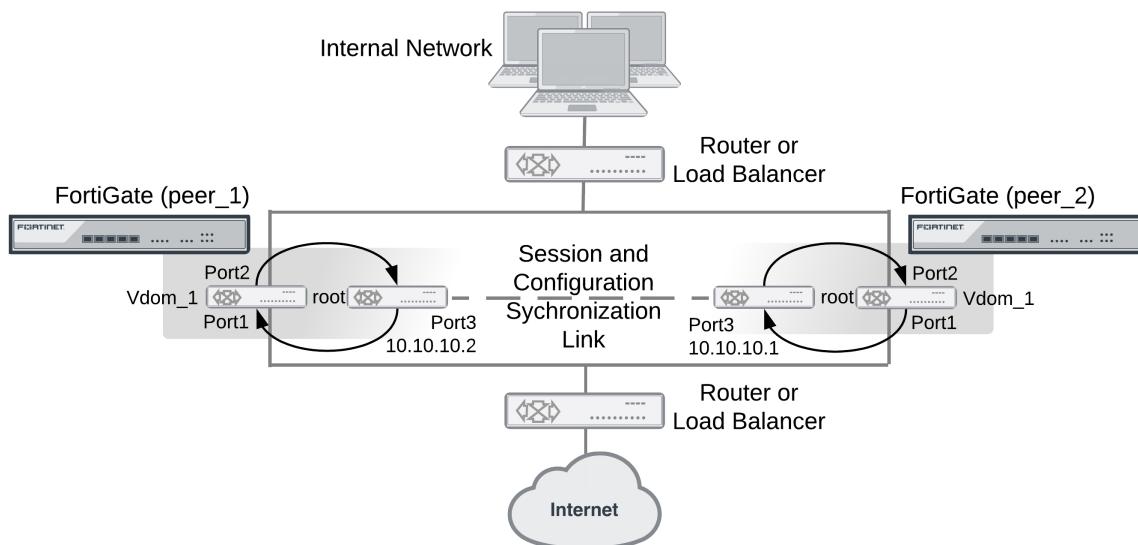
Basic example configuration

The following configuration example shows how to configure basic FGSP HA for the two peer FortiGates shown below.

- The host names of peers are peer_1 and peer_2.
- Both peers are configured with two virtual domains: root and vdom_1.
- All sessions processed by vdom_1 are synchronized.
- The synchronization link interface is port3 which is in the root virtual domain.
- The IP address of port3 on peer_1 is 10.10.10.1.
- The IP address of port3 on peer_2 is 10.10.10.2.

Also on both peers, port1 and port2 are added to vdom_1. On peer_1 the IP address of port1 is set to 192.168.20.1 and the IP address of port2 is set to 172.110.20.1. On peer_2 the IP address of port1 is set to 192.168.20.2 and the IP address of port2 is set to 172.110.20.2.

Example FGSP HA network configuration



To configure FGSP HA

1. Configure the load balancer or router to send all sessions to peer_1.
2. Configure the load balancer or router to send all traffic to peer_2 if peer_1 fails.
3. Use normal FortiGate configuration steps on peer_1:
 - Enable virtual domain configuration.
 - Add the vdom_1 virtual domain.
 - Add port1 and port2 to the vdom_1 virtual domain and configure these interfaces.
 - Set the IP address of port1 to 192.168.20.1.
 - Set the IP address of port2 to 172.110.20.1.

- Set the IP address of port3 to 10.10.10.1.
- Add route mode security policies between port1 and port2 to vdom_1.

4. Enter the following commands to configure session synchronization for peer_1:

```
config system cluster-sync
edit 1
set peerip 10.10.10.2
set peervd root
set syncvd vdom_1
end
```

5. Use normal FortiGate configuration steps on peer_2:

- Enable virtual domain configuration.
- Add the vdom_1 virtual domain.
- Add port1 and port2 to the vdom_1 virtual domain and configure these interfaces.
- Set the IP address of port1 to 192.168.20.2.
- Set the IP address of port2 to 172.110.20.2.
- Set the IP address of port3 to 10.10.10.1.
- Add route mode security policies between port1 and port2 to vdom_1.

6. Enter the following command to configure session synchronization for peer_1

```
config system cluster-sync
edit 1
set peerip 10.10.10.1
set peervd root
set syncvd vdom_1
end
```

Now that the FortiGates are connected and configured their configurations are synchronized, so when you make a configuration change on one FortiGate it is synchronized to the other one.

To add filters

You can add a filter to this basic configuration if you only want to synchronize some TCP sessions. For example you can enter the following command to add a filter so that only HTTP sessions are synchronized:

```
config system cluster-sync
edit 1
config filter
set service HTTP
end
end
```

You can also add a filter to control the source and destination addresses of the IPv4 packets that are synchronized. For example, you can enter the following to add a filter so that only sessions with source addresses in the range 10.10.10.100 to 10.10.10.200 are synchronized.

```
config system cluster-sync
edit 1
config filter
set srcaddr 10.10.10.100 10.10.10.200
end
end
```

You can also add a filter to control the source and destination addresses of the IPv6 packets that are synchronized. For example, you can enter the following to add a filter so that only sessions with destination addresses in the range 2001:db8:0:2::/64 are synchronized.

```
config system cluster-sync
  edit 1
    config filter
      set dstaddr6 2001:db8:0:2::/64
    end
  end
end
```

To synchronize TCP sessions

You enter the following to synchronization TCP sessions and set the synchronization link (heartbeat device):

```
config system ha
  set hbdev "port3" 50
  set session-pickup enable
end
```

To synchronize UDP and ICMP sessions

You enter the following to add synchronization of UDP and ICMP sessions to this configuration:

```
config system ha
  set session-pickup enable
  set session-pickup-connectionless enable
end
```

To synchronize the configuration

Enter the following to enable configuration synchronization.

```
config system ha
  set standalone-config-sync enable
end
```

Verifying the FGSP configuration and synchronization

You can use the following diagnose commands to verify that the FGSP and its synchronization functions are operating correctly.

FGSP configuration summary and status

Enter the following command to display a summary of the FGSP configuration and synchronization status:

```
diagnose sys session sync
sync_ctx: sync_started=1, sync_tcp=1, sync_others=1,
sync_expectation=1, sync_redir=0, sync_nat=1, stdalone_sessync=0.
sync: create=12:0, update=0, delete=0:0, query=14
recv: create=14:0, update=0, delete=0:0, query=12
ses pkts: send=0, alloc_fail=0, recv=0, recv_err=0 sz_err=0
nCfG_sess_sync_num=5, mtu=16000
sync_filter:
1: vd=0, szone=0, dzone=0, saddr=0.0.0.0:0.0.0.0, daddr=0.0.0.0:0.0.0.0,
```

`sync_started=1` shows that synchronization is working. If this is set to 0 then something is not correct with session synchronization and synchronization has not been able to start because of it.

`sync_tcp=1, sync_others=1, sync_expectation=1, and sync_nat=1` show that the FGSP has been configured to synchronize TCP, connectionless, asymmetric, and NAT sessions.

`sync: create=12:0 and recv: create=14:0` show that this FortiGate has synchronized 12 sessions to its peer and has received 14 sessions from its peer.

`sync_filter` shows the configured FGSP filter. In this case no filter has been created so all sessions are synchronized.

`vd=0` indicates that root VDOM sessions are synchronized.

Verifying that sessions are synchronized

Enter the command `diagnose sys session list` to display information about the sessions being processed by the FortiGate. In the command output look for sessions that should be synchronized and make sure they contain output lines that include `syncd` for example, `state=log may_dirty ndr syncd`) to confirm that they are being synchronized by the FGSP.

```
diagnose sys session list
session info: proto=6 proto_state=05 duration=469 expire=0 timeout=3600
flags=00000000 sockflag=00000000 sockport=21 av_idx=0 use=4
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 policy_dir=0 tunnel=/
state=log may_dirty ndr syncd
statistic(bytes/packets/allow_err): org=544/9/1 reply=621/7/0 tuples=2
origin->sink: org pre->post, reply pre->post dev=46->45/45->46
gwy=10.2.2.1/10.1.1.1
hook=pre dir=org act=noop 192.168.1.50:45327->172.16.1.100:21(0.0.0.0:0)
hook=post dir=reply act=noop 172.16.1.100:21->192.168.1.50:45327(0.0.0.0:0)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=1 id_policy_id=0 auth_info=0 chk_client_info=0 vd=0
serial=00002deb tos=ff/ff ips_view=1 app_list=2000 app=16427
dd_type=0 dd_mode=0
per_ip_bandwidth meter: addr=192.168.1.50, bps=633
```

Chapter 13 - IPsec VPN

Introduction

This FortiOS Handbook chapter contains the following sections:

[IPsec VPN concepts](#) explains the basic concepts that you need to understand about virtual private networks (VPNs).

[IPsec VPN overview](#) provides a brief overview of IPsec technology and includes general information about how to configure IPsec VPNs using this guide.

[IPsec VPN in the web-based manager](#) describes the IPsec VPN menu of the web-based manager interface.

[Gateway-to-gateway configurations](#) explains how to set up a basic gateway-to-gateway (site-to-site) IPsec VPN. In a gateway-to-gateway configuration, two FortiGate units create a VPN tunnel between two separate private networks.

[Hub-and-spoke configurations](#) describes how to set up hub-and-spoke IPsec VPNs. In a hub-and-spoke configuration, connections to a number of remote peers and/or clients radiate from a single, central FortiGate hub.

[Dynamic DNS configuration](#) describes how to configure a site-to-site VPN, in which one FortiGate unit has a static IP address and the other FortiGate unit has a dynamic IP address and a domain name.

[FortiClient dialup-client configurations](#) guides you through configuring a FortiClient dialup-client IPsec VPN. In a FortiClient dialup-client configuration, the FortiGate unit acts as a dialup server and VPN client functionality is provided by the FortiClient Endpoint Security application installed on a remote host.

[FortiGate dialup-client configurations](#) explains how to set up a FortiGate dialup-client IPsec VPN. In a FortiGate dialup-client configuration, a FortiGate unit with a static IP address acts as a dialup server and a FortiGate unit with a dynamic IP address initiates a VPN tunnel with the FortiGate dialup server.

[Supporting IKE Mode config clients](#) explains how to set up a FortiGate unit as either an IKE Mode Config server or client. IKE Mode Config is an alternative to DHCP over IPsec.

[Internet-browsing configuration](#) explains how to support secure web browsing performed by dialup VPN clients, and hosts behind a remote VPN peer. Remote users can access the private network behind the local FortiGate unit and browse the Internet securely. All traffic generated remotely is subject to the security policy that controls traffic on the private network behind the local FortiGate unit.

[Redundant VPN configurations](#) discusses the options for supporting redundant and partially redundant tunnels in an IPsec VPN configuration. A FortiGate unit can be configured to support redundant tunnels to the same remote peer if the FortiGate unit has more than one interface to the Internet.

[Transparent mode VPNs](#) describes two FortiGate units that create a VPN tunnel between two separate private networks transparently. In transparent mode, all FortiGate unit interfaces except the management interface are invisible at the network layer.

[IPv6 IPsec VPNs](#) describes FortiGate unit VPN capabilities for networks based on IPv6 addressing. This includes IPv4-over-IPv6 and IPv6-over-IPv4 tunnelling configurations. IPv6 IPsec VPNs are available in FortiOS 3.0 MR5 and later.

[L2TP and IPsec \(Microsoft VPN\)](#) explains how to support Microsoft Windows native VPN clients.

[GRE over IPsec \(Cisco VPN\)](#) explains how to interoperate with Cisco VPNs that use Generic Routing Encapsulation (GRE) protocol with IPsec.

[Protecting OSPF with IPsec](#) provides an example of protecting OSPF links with IPsec.

[Redundant OSPF routing over IPsec](#) provides an example of redundant secure communication between two remote networks using an OSPF VPN connection.

[OSPF over dynamic IPsec](#) provides an example of how to create a dynamic IPsec VPN tunnel that allows OSPF.

[BGP over dynamic IPsec](#) provides an example of how to create a dynamic IPsec VPN tunnel that allows BGP.

[Phase 1 parameters](#) provides detailed step-by-step procedures for configuring a FortiGate unit to accept a connection from a remote peer or dialup client. The basic Phase 1 parameters identify the remote peer or clients and support authentication through preshared keys or digital certificates. You can increase VPN connection security further using methods such as extended authentication (XAuth).

[Phase 2 parameters](#) provides detailed step-by-step procedures for configuring an IPsec VPN tunnel. During Phase 2, the specific IPsec security associations needed to implement security services are selected and a tunnel is established.

[Defining VPN security policies](#) explains how to specify the source and destination IP addresses of traffic transmitted through an IPsec VPN tunnel, and how to define a security encryption policy. Security policies control all IP traffic passing between a source address and a destination address.

[Logging and monitoring](#) and [Troubleshooting](#) provide VPN monitoring and troubleshooting procedures.

What's new in FortiOS 6.0.2

The following list contains new IPsec VPN features added in FortiOS 6.0.2. Click on a link to navigate to that section for further information.

- ["OCVPN support for High Availability \(HA\)" on page 1719](#)

What's new in FortiOS 6.0.1

The following list contains new IPsec VPN features added in FortiOS 6.0.1. Click on a link to navigate to that section for further information.

- [Updates to "IPsec support for ChaCha20/Poly1305 AEAD cipher" on page 1680](#)
- ["IPsec support for AES-GCM for IKEv2 Phase 1" on page 1680](#)

What's new in FortiOS 6.0

The following list contains new IPsec VPN features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- [Updates to "IPsec VPN Wizard options" on page 1625](#)
- ["Curve25519 128-bit elliptic curve group" on page 1630](#)
- ["Full CA chain checking" on page 1632](#)
- ["Timeout field in IPsec Monitor page" on page 1654](#)
- ["Dead Peer Detection" on page 1669](#)
- ["Quantum resistant IKEv2 SA negotiation" on page 1671](#)
- ["IPsec support for ChaCha20/Poly1305 AEAD cipher" on page 1680](#)
- ["Remote Internet browsing for Site-to-Site VPN from the IPsec VPN Wizard" on page 1694](#)
- ["One-Click VPN \(OCVPN\)" on page 1719](#)
- ["Split-exclude in IKEv1 mode-cfg" on page 1758](#)
- ["IPsec VPN tunnel aggregate interfaces" on page 1767](#)
- ["Changing GRE over GRE tunnel interface attributes" on page 1801](#)
- ["IPv6 support for GRE tunnels" on page 1801](#)

IPsec VPN concepts

Virtual Private Network (VPN) technology enables remote users to connect to private computer networks to gain access to their resources in a secure way. For example, an employee traveling or working from home can use a VPN to securely access the office network through the Internet.

Instead of remotely logging on to a private network using an unencrypted and unsecure Internet connection, the use of a VPN ensures that unauthorized parties cannot access the office network and cannot intercept any of the information that is exchanged between the employee and the office. It is also common to use a VPN to connect the private networks of two or more offices.

Fortinet offers VPN capabilities in the FortiGate Unified Threat Management (UTM) appliance and in the FortiClient Endpoint Security suite of applications. A FortiGate unit can be installed on a private network, and FortiClient software can be installed on the user's computer. It is also possible to use a FortiGate unit to connect to the private network instead of using FortiClient software.


This chapter discusses VPN terms and concepts including:

VPN tunnels

The data path between a user's computer and a private network through a VPN is referred to as a tunnel. Like a physical tunnel, the data path is accessible only at both ends. In the telecommuting scenario, the tunnel runs between the FortiClient application on the user's PC, or a FortiGate unit or other network device and the FortiGate unit on the office private network.

Encapsulation makes this possible. IPsec packets pass from one end of the tunnel to the other and contain data packets that are exchanged between the local user and the remote private network. Encryption of the data packets ensures that any third-party who intercepts the IPsec packets can not access the data.

Encoded data going through a VPN tunnel



You can create a VPN tunnel between:

- A PC equipped with the FortiClient application and a FortiGate unit
- Two FortiGate units
- Third-party VPN software and a FortiGate unit

For more information on third-party VPN software, refer to the [Fortinet Knowledge Base](#) for more information.

Tunnel templates

Several tunnel templates are available in the IPsec VPN Wizard that cover a variety of different types of IPsec VPN. A list of these templates appear on the first page of the Wizard, located at **VPN > IPsec Wizard**. The tunnel template list follows.

IPsec VPN Wizard options

VPN Type	Remote Device Type		NAT Options	Description
Site to Site	FortiGate		<ul style="list-style-type: none"> No NAT between sites This site is behind NAT The remote site is behind NAT 	Static tunnel between this FortiGate and a remote FortiGate.
	Cisco		<ul style="list-style-type: none"> No NAT between sites This site is behind NAT The remote site is behind NAT 	Static tunnel between this FortiGate and a remote Cisco firewall.
Remote Access	Client-based	FortiClient VPN for OS X, Windows, and Android	N/A	On-demand tunnel for users using the FortiClient software.
		Cisco AnyConnect	N/A	On-demand tunnel for users using the Cisco IPsec client.
	Native	iOS Native	N/A	On-demand tunnel for iPhone/iPad users using the native iOS IPsec client.
		Android Native	N/A	On-demand tunnel for Android users using the native L2TP/IPsec client.
		Windows Native	N/A	On-demand tunnel for Android users using the native L2TP/IPsec client.
Custom	N/A		N/A	No Template.



Cisco's VPN Client has reached its End-of-Life/End-of-Support as of July 30, 2016, and has been replaced by [Cisco AnyConnect Secure Mobility Client](#).



In FortiOS 5.6.4+, the first step of the VPN Creation Wizard (**VPN > IPsec Wizard**) delineates the **Remote Device Type** (for **Remote Access** templates) between **Client-based** and **Native** in order to distinguish FortiClient and Cisco device options from native OS device options.

VPN tunnel list

Once you create an IPsec VPN tunnel, it appears in the VPN tunnel list at **VPN > IPsec Tunnels**. By default, the tunnel list indicates the name of the tunnel, its interface binding, the tunnel template used, and the tunnel status. If you right-click on the table header row, you can include columns for comments, IKE version, mode (aggressive vs main), phase 2 proposals, and reference number. The tunnel list page also includes the option to create a new tunnel, as well as the options to edit or delete a highlighted tunnel.

FortiView VPN tunnel map

A geospatial map can be found under **FortiView > VPN Map** to help visualize IPsec (and SSL) VPN connections to a FortiGate using Google Maps. This feature adds a geographical-IP API service for resolving spatial locations from IP addresses.


VPN gateways

A gateway is a router that connects the local network to other networks. The default gateway setting in your computer's TCP/IP properties specifies the gateway for your local network.

A VPN gateway functions as one end of a VPN tunnel. It receives incoming IPsec packets, decrypts the encapsulated data packets and passes the data packets to the local network. Also, it encrypts data packets destined for the other end of the VPN tunnel, encapsulates them, and sends the IPsec packets to the other VPN gateway. The VPN gateway is a FortiGate unit because the private network behind it is protected, ensuring the security of the unencrypted VPN data. The gateway can also be FortiClient software running on a PC since the unencrypted data is secure on the PC.

The IP address of a VPN gateway is usually the IP address of the network interface that connects to the Internet. Optionally, you can define a secondary IP address for the interface and use that address as the local VPN gateway address. The benefit of doing this is that your existing setup is not affected by the VPN settings.

The following diagram shows a VPN connection between two private networks with FortiGate units acting as the VPN gateways. This configuration is commonly referred to as Gateway-to-Gateway IPsec VPN.

VPN tunnel between two private networks

Although the IPsec traffic may actually pass through many Internet routers, you can visualize the VPN tunnel as a simple secure connection between the two FortiGate units.

Users on the two private networks do not need to be aware of the VPN tunnel. The applications on their computers generate packets with the appropriate source and destination addresses, as they normally do. The FortiGate units manage all the details of encrypting, encapsulating, and sending the packets to the remote VPN gateway.

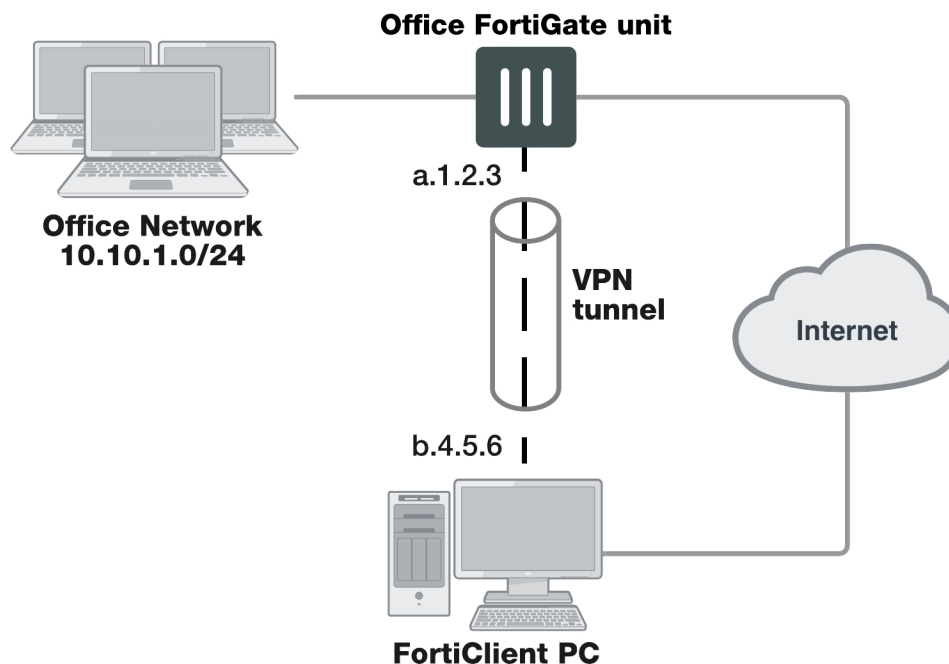
The data is encapsulated in IPsec packets only in the VPN tunnel between the two VPN gateways. Between the user's computer and the gateway, the data is on the secure private network and it is in regular IP packets.

For example User1 on the Site A network, at IP address 10.10.1.7, sends packets with destination IP address 192.168.10.8, the address of User2 on the Site B network. The Site A FortiGate unit is configured to send packets with destinations on the 192.168.10.0 network through the VPN, encrypted and encapsulated. Similarly, the Site B FortiGate unit is configured to send packets with destinations on the 10.10.1.0 network through the VPN tunnel to the Site A VPN gateway.

In the site-to-site, or gateway-to-gateway VPN shown below, the FortiGate units have static (fixed) IP addresses and either unit can initiate communication.

You can also create a VPN tunnel between an individual PC running FortiClient and a FortiGate unit, as shown below. This is commonly referred to as Client-to-Gateway IPsec VPN.

VPN tunnel between a FortiClient PC and a FortiGate unit



On the PC, the FortiClient application acts as the local VPN gateway. Packets destined for the office network are encrypted, encapsulated into IPsec packets, and sent through the VPN tunnel to the FortiGate unit. Packets for other destinations are routed to the Internet as usual. IPsec packets arriving through the tunnel are decrypted to recover the original IP packets.

Clients, servers, and peers

A FortiGate unit in a VPN can have one of the following roles:

- **Server** — responds to a request to establish a VPN tunnel.
- **Client** — contacts a remote VPN gateway and requests a VPN tunnel.
- **Peer** — brings up a VPN tunnel or responds to a request to do so.

The site-to-site VPN shown above is a peer-to-peer relationship. Either FortiGate unit VPN gateway can establish the tunnel and initiate communications. The FortiClient-to-FortiGate VPN shown below is a client-server relationship. The FortiGate unit establishes a tunnel when the FortiClient PC requests one.

A FortiGate unit cannot be a VPN server if it has a dynamically-assigned IP address. VPN clients need to be configured with a static IP address for the server. A FortiGate unit acts as a server only when the remote VPN gateway has a dynamic IP address or is a client-only device or application, such as FortiClient.

As a VPN server, a FortiGate unit can also offer automatic configuration for FortiClient PCs. The user needs to know only the IP address of the FortiGate VPN server and a valid user name/password. FortiClient downloads the VPN configuration settings from the FortiGate VPN server. For information about configuring a FortiGate unit as a VPN server, see the [FortiClient Administration Guide](#).

Encryption

Encryption mathematically transforms data to appear as meaningless random numbers. The original data is called plaintext and the encrypted data is called ciphertext. The opposite process, called decryption, performs the inverse operation to recover the original plaintext from the ciphertext.

The process by which the plaintext is transformed to ciphertext and back again is called an algorithm. All algorithms use a small piece of information, a key, in the arithmetic process of converted plaintext to ciphertext, or vice-versa. IPsec uses symmetrical algorithms, in which the same key is used to both encrypt and decrypt the data. The security of an encryption algorithm is determined by the length of the key that it uses. FortiGate IPsec VPNs offer the following encryption algorithms, in descending order of security:

Encryption	Description
ChaCha20/Poly1305	A combination of the ChaCha20 symmetric cipher and Poly1305-AES, a variant of the AES 128-bit block algorithm that uses a 128-bit key and an 128-bit nonce.
AES-GCM	Galois/Counter Mode (GCM), a block cipher mode of operation providing both confidentiality and data origin authentication.
AES256	A 128-bit block algorithm that uses a 256-bit key.
AES192	A 128-bit block algorithm that uses a 192-bit key.
AES128	A 128-bit block algorithm that uses a 128-bit key.
3DES	Triple-DES, in which plain text is DES-encrypted three times by three keys.
DES	Digital Encryption Standard, a 64-bit block algorithm that uses a 56-bit key

The default encryption algorithms provided on FortiGate units make recovery of encrypted data almost impossible without the proper encryption keys.

There is a human factor in the security of encryption. The key must be kept secret, known only to the sender and receiver of the messages. Also, the key must not be something that unauthorized parties might easily guess, such as the sender's name, birthday or simple sequence such as 123456.

Diffie-Hellman groups

FortiOS IPsec VPN supports the following Diffie-Hellman (DH) asymmetric key algorithms for public key cryptography.

DH Group	Description
1	More Modular Exponential (MODP) DH Group with a 768-bit modulus
2	MODP with a 1024-bit modulus
5	MODP with a 1536-bit modulus

DH Group	Description
14	MODP with a 2048-bit modulus
15	MODP with a 3027-bit modulus
16	MODP with a 4096-bit modulus
17	MODP with a 6144-bit modulus
18	MODP with a 8192-bit modulus
19	256-bit random elliptic curve group
20	384-bit random elliptic curve group
21	521-bit random elliptic curve group
27	Brainpool 224-bit elliptic curve group
28	Brainpool 256-bit elliptic curve group
29	Brainpool 384-bit elliptic curve group
30	Brainpool 512-bit elliptic curve group
31	Curve25519 128-bit elliptic curve group

* When using aggressive mode, DH groups cannot be negotiated.

By default, DH group 14 is selected, to provide sufficient protection for stronger cipher suites that include AES and SHA2. If you select multiple DH groups, the order they appear in the configuration is the order in which they are negotiated.

If both VPN peers (or a VPN server and its client) have static IP addresses and use aggressive mode, select a single DH group. The setting on the FortiGate unit must be identical to the setting on the remote peer or dialup client.

When the remote VPN peer or client has a dynamic IP address and uses aggressive mode, select up to three DH groups on the FortiGate unit and one DH group on the remote peer or dialup client. The setting on the remote peer or dialup client must be identical to one of the selections on the FortiGate unit.

If the VPN peer or client employs main mode, you can select multiple DH groups. At least one of the settings on the remote peer or dialup client must be identical to the selections on the FortiGate unit.

IPsec overheads

The FortiGate sets an IPsec tunnel Maximum Transmission Unit (MTU) of 1436 for 3DES/SHA1 and an MTU of 1412 for AES128/SHA1, as seen with `diag vpn tunnel list`. This indicates that the FortiGate allocates 64 bytes of overhead for 3DES/SHA1 and 88 bytes for AES128/SHA1, which is the difference if you subtract this MTU from a typical ethernet MTU of 1500 bytes.

During the encryption process, AES/DES operates using a specific size of data which is block size. If data is smaller than that, it will be padded for the operation. MD5/SHA-1 HMAC also operates using a specific block size.

The following table describes the potential maximum overhead for each IPsec encryption:

IPsec Transform Set	IPsec Overhead (Max. bytes)
ESP-AES (256, 192, or 128), ESP-SHA-HMAC, or MD5	88
ESP-AES (256, 192, or 128)	61
ESP-3DES, ESP-DES	45
ESP-(DES or 3DES), ESP-SHA-HMAC, or MD5	64
ESP-Null, ESP-SHA-HMAC, or MD5	45
AH-SHA-HMAC or MD5	44

Authentication

To protect data via encryption, a VPN must ensure that only authorized users can access the private network. You must use either a preshared key on both VPN gateways or RSA X.509 security certificates. The examples in this guide use only preshared key authentication. Refer to the [Fortinet Knowledge Base](#) for articles on RSA X.509 security certificates.

Preshared keys

A preshared key contains at least six random alphanumeric characters. Users of the VPN must obtain the preshared key from the person who manages the VPN server and add the preshared key to their VPN client configuration.

Although it looks like a password, the preshared key, also known as a shared secret, is never sent by either gateway. The preshared key is used in the calculations at each end that generate the encryption keys. As soon as the VPN peers attempt to exchange encrypted data, preshared keys that do not match will cause the process to fail.

Additional authentication

To increase security, you can require additional means of authentication from users, such as:

- An identifier, called a peer ID or a local ID.
- Extended authentication (XAUTH) which imposes an additional user name/password requirement.

A Local ID is an alphanumeric value assigned in the Phase 1 configuration. The Local ID of a peer is called a Peer ID.

In FortiOS 5.2, new authentication methods have been implemented for IKE: ECDSA-256, ECDSA-384, and ECDSA-521. However, AES-XCBC is not supported.

Full CA chain checking

Added a new option (enabled by default) to fail certificate verification if any of the CAs in the trust chain are not found in the CA store. When disabled, a sub-CA is sufficient to pass certificate verification.

Syntax

```
config vpn certificate setting
    set check-ca-chain {enable | disable}
end
```

Phase 1 and Phase 2 settings

A VPN tunnel is established in two phases: Phase 1 and Phase 2. Several parameters determine how this is done. Except for IP addresses, the settings simply need to match at both VPN gateways. There are defaults that are appropriate for most cases.

FortiClient distinguishes between Phase 1 and Phase 2 only in the VPN Advanced settings and uses different terms. Phase 1 is called the IKE Policy. Phase 2 is called the IPsec Policy.

Phase 1

In Phase 1, the two VPN gateways exchange information about the encryption algorithms that they support and then establish a temporary secure connection to exchange authentication information.

When you configure your FortiGate unit or FortiClient application, you must specify the following settings for Phase 1:

Remote gateway	The remote VPN gateway's address. FortiGate units also have the option of operating only as a server by selecting the "Dialup User" option.
Preshared key	This must be the same at both ends. It is used to encrypt Phase 1 authentication information.
Local interface	The network interface that connects to the other VPN gateway. This applies on a FortiGate unit only.

All other Phase 1 settings have default values. These settings mainly configure the types of encryption to be used. The default settings on FortiGate units and in the FortiClient application are compatible. The examples in this guide use these defaults.

For more detailed information about Phase 1 settings, see [Phase 1 parameters on page 1655](#).

Phase 2

Similar to the Phase 1 process, the two VPN gateways exchange information about the encryption algorithms that they support for Phase 2. You may choose different encryption for Phase 1 and Phase 2. If both gateways have at least one encryption algorithm in common, a VPN tunnel can be established. Keep in mind that more algorithms each phase does not share with the other gateway, the longer negotiations will take. In extreme cases this may cause timeouts during negotiations.

To configure default Phase 2 settings on a FortiGate unit, you need only select the name of the corresponding Phase 1 configuration. In FortiClient, no action is required to enable default Phase 2 settings.

For more detailed information about Phase 2 settings, see [Phase 2 parameters on page 1675](#).

Security Association

The establishment of a Security Association (SA) is the successful outcome of Phase 1 negotiations. Each peer maintains a database of information about VPN connections. The information in each SA can include cryptographic algorithms and keys, keylife, and the current packet sequence number. This information is kept synchronized as the VPN operates. Each SA has a Security Parameter Index (SPI) that is provided to the remote peer at the time the SA is established. Subsequent IPsec packets from the peer always reference the relevant SPI. It is possible for peers to have multiple VPNs active simultaneously, and correspondingly multiple SPIs.

The IPsec SA connect message generated is used to install dynamic selectors. These selectors can be installed via the auto-negotiate mechanism. When phase 2 has auto-negotiate enabled, and phase 1 has mesh selector-type set to **subnet**, a new dynamic selector will be installed for each combination of source and destination subnets. Each dynamic selector will inherit the auto-negotiate option from the template selector and begin SA negotiation. Phase 2 selector sources from dial-up clients will all establish SAs without traffic being initiated from the client subnets to the hub.

Remote IP address change detection

SAs are stored in a hash table when keyed off the IPsec SA SPI value. This enables the FortiGate, for each inbound ESP packet received, to immediately look up the SA and compare the stored IP address against the one in the incoming packet. If the incoming and stored IP addresses differ, an IP address change can be made in the kernel SA, and an update event can be triggered for IKE.

IKE and IPsec packet processing

Internet Key Exchange (IKE) is the protocol used to set up SAs in IPsec negotiation. As described in [Phase 1 parameters on page 1655](#), you can optionally choose IKEv2 over IKEv1 if you configure a route-based IPsec VPN. IKEv2 simplifies the negotiation process, in that it provides no choice of Aggressive or Main mode in Phase 1. IKEv2 also uses less bandwidth.

The following sections identify how IKE versions 1 and 2 operate and differentiate.

IKEv1

Phase 1

A peer, identified in the IPsec policy configuration, begins the IKE negotiation process. This IKE Security Association (SA) agreement is known as Phase 1. The Phase 1 parameters identify the remote peer or clients and supports authentication through pre-shared key (PSK) or digital certificate. You can increase access security further using peer identifiers, certificate distinguished names, group names, or the FortiGate extended authentication (XAuth) option for authentication purposes. Basically, Phase 1 authenticates a remote peer and sets up a secure communication channel for establishing Phase 2, which negotiates the IPsec SA.

IKE Phase 1 can occur in either Main mode or Aggressive mode. For more information, see [Phase 1 parameters on page 1655](#).

IKE Phase 1 is successful only when the following are true:

- Each peer negotiates a matching IKE SA policy.
- Each peer is authenticated and their identities protected.
- The Diffie-Hellman exchange is authenticated (the pre-shared secret keys match).

For more information on Phase 1, see [Phase 1 parameters on page 1655](#).

Phase 2

Phase 2 parameters define the algorithms that the FortiGate unit can use to encrypt and transfer data for the remainder of the session in an IPsec SA. The basic Phase 2 settings associate IPsec Phase 2 parameters with a Phase 1 configuration.

In Phase 2, the VPN peer or client and the FortiGate unit exchange keys again to establish a more secure communication channel. The Phase 2 Proposal parameters select the encryption and authentication algorithms needed to generate keys for protecting the implementation details of the SA. The keys are generated automatically using a Diffie-Hellman algorithm.

In Phase 2, Quick mode selectors determine which IP addresses can perform IKE negotiations to establish a tunnel. By only allowing authorized IP addresses access to the VPN tunnel, the network is more secure. For more information, see [Phase 2 parameters on page 1675](#).

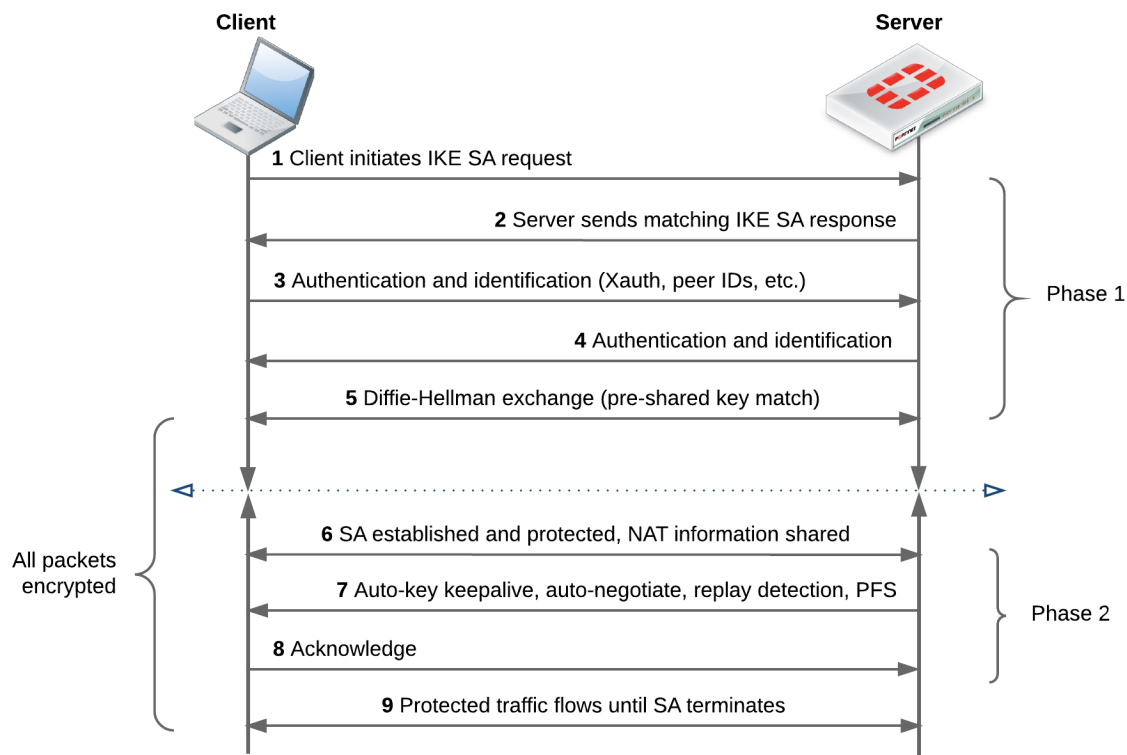
IKE Phase 2 is successful only when the following are true:

- The IPsec SA is established and protected by the IKE SA.
- The IPsec SA is configured to renegotiate after set durations (see [Phase 2 parameters on page 1675](#) and [Phase 2 parameters on page 1675](#)).
- **Optional:** Replay Detection is enabled. Replay attacks occur when an unauthorized party intercepts a series of IPsec packets and replays them back into the tunnel. See [Phase 2 parameters on page 1675](#).
- **Optional:** Perfect Forward Secrecy (PFS) is enabled. PFS improves security by forcing a new Diffie-Hellman exchange whenever keylife expires. See [Phase 2 parameters on page 1675](#).

For more information on Phase 2, see [Phase 2 parameters on page 1675](#).

With Phase 2 established, the IPsec tunnel is fully negotiated and traffic between the peers is allowed until the SA terminates (for any number of reasons; time-out, interruption, disconnection, etc).

The entire IKEv1 process is demonstrated in the following diagram:



IKEv2

Phase 1

Unlike Phase 1 of IKEv1, IKEv2 does not provide options for Aggressive or Main mode. Furthermore, Phase 1 of IKEv2 begins immediately with an IKE SA initiation, consisting of only two packets (containing all the information typically contained in four packets for IKEv1), securing the channel such that all following transactions are encrypted (see [Phase 1 parameters on page 1655](#)).

The encrypted transactions contain the IKE authentication, since remote peers have yet to be authenticated. This stage of IKE authentication in IKEv2 can loosely be called Phase 1.5.

Phase 1.5

As part of this phase, IKE authentication must occur. IKE authentication consists of the following:

- The authentication payloads and Internet Security Association and Key Management Protocol (ISAKMP) identifier.
- The authentication method (RSA, PSK, ECDSA, or EAP).
- The IPsec SA parameters.

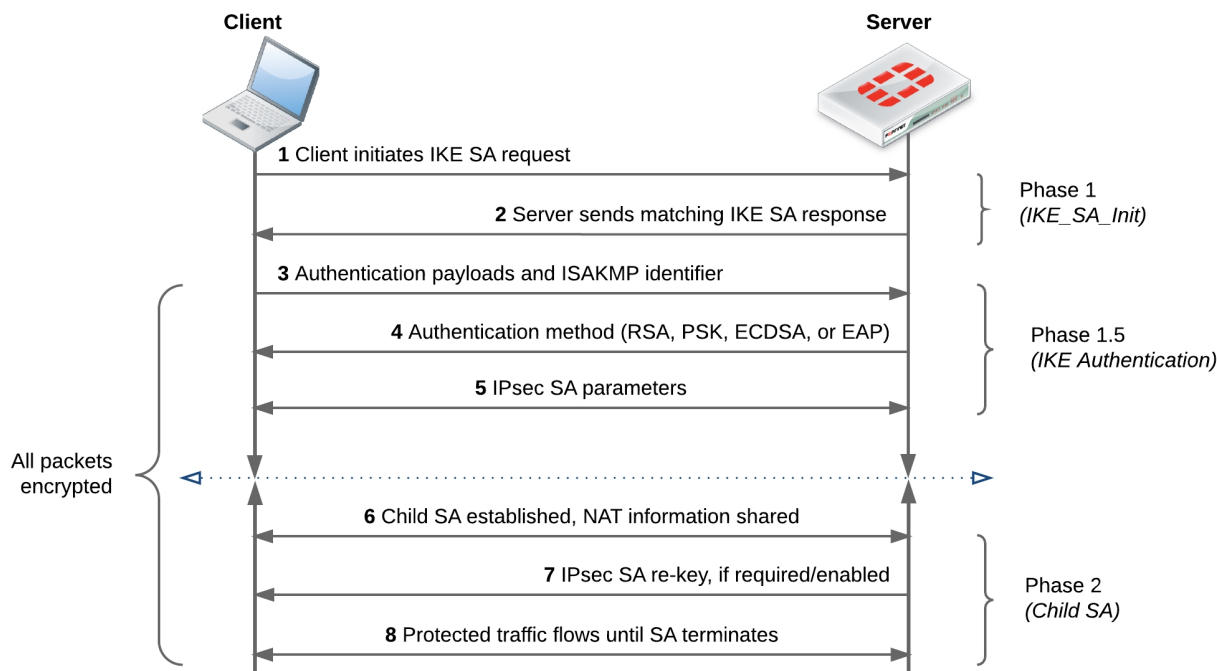
Due to the number of authentication methods potentially used, and SAs established, the overall IKEv2 negotiation can range from 4 packets (no EAP exchange at all) to many more.

At this point, both peers have a security association complete and ready to encrypt traffic.

Phase 2

In IKEv1, Phase 2 uses Quick mode to negotiate an IPsec SA between peers. In IKEv2, since the IPsec SA is already established, Phase 2 is essentially only used to negotiate “child” SAs, or to re-key an IPsec SA. That said, there are only two packets for each exchange of this type, similar to the exchange at the outset of Phase 1.5.

The entire IKEv2 process is demonstrated in the following diagram:



Support for IKEv2 session resumption

If a gateway loses connectivity to the network, clients can attempt to re-establish the lost session by presenting the ticket to the gateway (as described in [RFC 5723](#)). As a result, sessions can be resumed much faster, as DH exchange that is necessary to establish a brand new connection is skipped. This feature implements "ticket-by-value", whereby all information necessary to restore the state of a particular IKE SA is stored in the ticket and sent to the client.

IKEv2 asymmetric authentication

Asymmetric authentication allows both sides of an authentication exchange to use different authentication methods, for example the initiator may be using a shared key, while the responder may have a public signature key and certificate.

The command `authmethod-remote` is available under `config vpn ipsec phase1-interface`.

For more detailed information on authentication of the IKE SA, see [RFC 5996 - Internet Key Exchange Protocol Version 2 \(IKEv2\)](#).

IKEv2 Digital Signature Authentication support

FortiOS supports the use of Digital Signature authentication, which changes the format of the Authentication Data payload in order to support different signature methods.

Instead of just containing a raw signature value calculated as defined in the original IKE RFCs, the Auth Data now includes an ASN.1 formatted object that provides details on how the signature was calculated, such as the signature type, hash algorithm, and signature padding method.

For more detailed information on IKEv2 Digital Signature authentication, see [RFC 7427 - Signature Authentication in the Internet Key Exchange Version 2 \(IKEv2\)](#).

Unique IKE identifiers

When enabled, the following `phase1` CLI command (`enforce-unique-id`) requires all IPsec VPN clients to use a unique identifier when connecting.

CLI syntax

```
config vpn ipsec phase1
  edit <name>
    set enforce-unique-id {keep-new | keep-old | disable} Default is disable.
  next
end
```

Use `keep-new` to replace the old connection if an ID collision is detected on the gateway.

Use `keep-old` to reject the new connection if an ID collision is detected.

IKEv2 ancillary RADIUS group authentication

This feature provides for the IDi information to be extracted from the IKEv2 AUTH exchange and sent to a RADIUS server, along with a fixed password (configurable via CLI only), to perform an additional group authentication step prior to tunnel establishment. The RADIUS server may return framed-IP-address, framed-ip-netmask, and dns-server attributes, which are then applied to the tunnel.

It should be noted, unlike Xauth or EAP, this feature does not perform individual user authentication, but rather treats all users on the gateway as a single group, and authenticates that group with RADIUS using a fixed password. This feature also works with RADIUS accounting, including the `phase1 acct-verify` option.

Syntax

```
config vpn ipsec phase1-interface
  edit <name>
    set mode-cfg enable
    set type dynamic
    set ike-version 2
    set group-authentication {enable | disable}
    set group-authentication-secret <password>
  next
end
```

IPsec VPN overview

This section provides a brief overview of IPsec technology and includes general information about how to configure IPsec VPNs using this guide.

The following topics are included in this section:

VPN configurations interact with the firewall component of the FortiGate unit. There must be a security policy in place to permit traffic to pass between the private network and the VPN tunnel.

Security policies for VPNs specify:

- The FortiGate interface that provides the physical connection to the remote VPN gateway, usually an interface connected to the Internet
- The FortiGate interface that connects to the private network
- IP addresses associated with data that has to be encrypted and decrypted
- Optionally, a schedule that restricts when the VPN can operate
- Optionally, the services (types of data) that can be sent

When the first packet of data that meets all of the conditions of the security policy arrives at the FortiGate unit, a VPN tunnel may be initiated and the encryption or decryption of data is performed automatically afterward. For more information, see [Defining VPN security policies on page 1](#).

Where possible, you should create route-based VPNs. Generally, route-based VPNs are more flexible and easier to configure than policy-based VPNs — by default they are treated as interfaces. However, these two VPN types have different requirements that limit where they can be used.

Types of VPNs

FortiGate unit VPNs can be policy-based or route-based. There is little difference between the two types. In both cases, you specify Phase 1 and Phase 2 settings. However there is a difference in implementation. A route-based VPN creates a virtual IPsec network interface that applies encryption or decryption as needed to any traffic that it carries. That is why route-based VPNs are also known as interface-based VPNs. A policy-based VPN is implemented through a special security policy that applies the encryption you specified in the Phase 1 and Phase 2 settings.

Route-based VPNs

For a route-based VPN, you create two security policies between the virtual IPsec interface and the interface that connects to the private network. In one policy, the virtual interface is the source. In the other policy, the virtual interface is the destination. This creates bidirectional policies that ensure traffic will flow in both directions over the VPN.

A route-based VPN is also known as an interface-based VPN.



Each route-based IPsec VPN tunnel requires a virtual IPsec interface. As such, the amount of possible route-based IPsec VPNs is limited by the **system.interface** table size. The **system.interface** table size for most devices is 8192.

For a complete list of table sizes for all devices, refer to the [Maximum Values](#) table.

Policy-based VPNs

For a policy-based VPN, one security policy enables communication in both directions. You enable inbound and outbound traffic as needed within that policy, or create multiple policies of this type to handle different types of traffic differently. For example HTTPS traffic may not require the same level of scanning as FTP traffic.

A policy-based VPN is also known as a tunnel-mode VPN.

Comparing policy-based or route-based VPNs

For both VPN types you create Phase 1 and Phase 2 configurations. Both types are handled in the stateful inspection security layer, assuming there is no IPS or AV. For more information on the three security layers, see the [FortiOS Troubleshooting guide](#).

The main difference is in the security policy.

You create a policy-based VPN by defining an IPSEC security policy between two network interfaces and associating it with the VPN tunnel (Phase 1) configuration.

You create a route-based VPN by creating a virtual IPsec interface. You then define a regular ACCEPT security policy to permit traffic to flow between the virtual IPsec interface and another network interface. And lastly, configure a static route to allow traffic over the VPN.

Where possible, you should create route-based VPNs. Generally, route-based VPNs are more flexible and easier to configure than policy-based VPNs — by default they are treated as interfaces. However, these two VPN types have different requirements that limit where they can be used.

Comparison of policy-based and route-based VPNs

Features	Policy-based	Route-based
Both NAT and transparent modes available	Yes	NAT mode only
L2TP-over-IPsec supported	Yes	Yes
GRE-over-IPsec supported	No	Yes
security policy requirements	Requires a security policy with IPSEC action that specifies the VPN tunnel	Requires only a simple security policy with ACCEPT action
Number of policies per VPN	One policy controls connections in both directions	A separate policy is required for connections in each direction

Planning your VPN

It is a good idea to plan the VPN configuration ahead of time. This will save time later and help you configure your VPN correctly.

All VPN configurations are comprised of numerous required and optional parameters. Before you begin, you need to determine:

- Where the IP traffic originates and where it needs to be delivered
- Which hosts, servers, or networks to include in the VPN
- Which VPN devices to include in the configuration
- Through which interfaces the VPN devices communicate
- Through which interfaces do private networks access the VPN gateways

Once you have this information, you can select a VPN topology that suits the network environment.

Network topologies

The topology of your network will determine how remote peers and clients connect to the VPN and how VPN traffic is routed.

VPN network topologies and brief descriptions

Topology	Description
Gateway-to-gateway configurations	Standard one-to-one VPN between two FortiGate units. See Gateway-to-gateway configurations on page 1 .
Hub-and-spoke configurations	One central FortiGate unit has multiple VPNs to other remote FortiGate units. See Hub-and-spoke configurations on page 1 .
Dynamic DNS configuration	One end of the VPN tunnel has a changing IP address and the other end must go to a dynamic DNS server for the current IP address before establishing a tunnel. See Dynamic DNS configuration on page 1 .
FortiClient dialup-client configurations	Typically remote FortiClient dialup-clients use dynamic IP addresses through NAT devices. The FortiGate unit acts as a dialup server allowing dialup VPN connections from multiple sources. See FortiClient dialup-client configurations on page 1 .
FortiGate dialup-client configurations	Similar to FortiClient dialup-client configurations but with more gateway-to-gateway settings such as unique user authentication for multiple users on a single VPN tunnel. See FortiGate dialup-client configurations on page 1 .
Internet-browsing configuration	Secure web browsing performed by dialup VPN clients, and/or hosts behind a remote VPN peer. See Internet-browsing configuration on page 1 .
Redundant VPN configurations	Options for supporting redundant and partially redundant IPsec VPNs, using route-based approaches. See Redundant VPN configurations on page 1 .
Transparent mode VPNs	In transparent mode, the FortiGate acts as a bridge with all incoming traffic being broadcast back out on all other interfaces. Routing and NAT must be performed on external routers. See Transparent mode VPNs on page 1 .
L2TP and IPsec (Microsoft VPN)	Configure VPN for Microsoft Windows dialup clients using the built in L2TP software. Users do not have to install any See L2TP and IPsec (Microsoft VPN) on page 1 .

These sections contain high-level configuration guidelines with cross-references to detailed configuration procedures. If you need more detail to complete a step, select the cross-reference in the step to drill-down to more detail. Return to the original procedure to complete the procedure. For a general overview of how to configure a VPN, see [Planning your VPN](#).

General preparation steps

A VPN configuration defines relationships between the VPN devices and the private hosts, servers, or networks making up the VPN. Configuring a VPN involves gathering and recording the following information. You will need this information to configure the VPN.

- **The private IP addresses of participating hosts, servers, and/or networks.** These IP addresses represent the source addresses of traffic that is permitted to pass through the VPN. A IP source address can be an individual IP address, an address range, or a subnet address.
- **The public IP addresses of the VPN end-point interfaces.** The VPN devices establish tunnels with each other through these interfaces.
- **The private IP addresses associated with the VPN-device interfaces to the private networks.** Computers on the private networks behind the VPN gateways will connect to their VPN gateways through these interfaces.

How to use this guide to configure an IPsec VPN

This guide uses a task-based approach to provide all of the procedures needed to create different types of VPN configurations. Follow the step-by-step configuration procedures in this guide to set up the VPN.

The following configuration procedures are common to all IPsec VPNs:

1. Define the Phase 1 parameters that the FortiGate unit needs to authenticate remote peers or clients and establish a secure connection. See [Phase 1 parameters on page 1655](#).
2. Define the Phase 2 parameters that the FortiGate unit needs to create a VPN tunnel with a remote peer or dialup client. See [Phase 2 parameters on page 1675](#).
3. Specify the source and destination addresses of IP packets that are to be transported through the VPN tunnel. See [Defining policy addresses on page 1](#).
4. Create an IPsec security policy to define the scope of permitted services between the IP source and destination addresses. See [Defining VPN security policies on page 1](#).



These steps assume you configure the FortiGate unit to generate unique IPsec encryption and authentication keys automatically. In situations where a remote VPN peer or client requires a specific IPsec encryption and authentication key, you must configure the FortiGate unit to use manual keys instead of performing Steps 1 and 2.

IPsec VPN in the web-based manager

To configure an IPsec VPN, use the general procedure below. With these steps, your FortiGate unit will automatically generate unique IPsec encryption and authentication keys. If a remote VPN peer or client requires a specific IPsec encryption or authentication key, you must configure your FortiGate unit to use manual keys instead.

1. Define Phase 1 parameters to authenticate remote peers and clients for a secure connection. See [IPsec VPN in the web-based manager on page 1642](#).
2. Define Phase 2 parameters to create a VPN tunnel with a remote peer or dialup client. See [IPsec VPN in the web-based manager on page 1642](#).
3. Create a security policy to permit communication between your private network and the VPN. Policy-based VPNs have an action of IPSEC, where for interface-based VPNs the security policy action is ACCEPT. See [Defining VPN security policies on page 1](#).

The FortiGate unit implements the Encapsulated Security Payload (ESP) protocol. Internet Key Exchange (IKE) is performed automatically based on pre-shared keys or X.509 digital certificates. Interface mode, supported in NAT mode only, creates a virtual interface for the local end of a VPN tunnel.

This chapter contains the following sections:

Phase 1 configuration

To begin defining the Phase 1 configuration, go to **VPN > IPsec Tunnels** and select **Create New**. Enter a unique descriptive name for the VPN tunnel and follow the instructions in the VPN Creation Wizard.

The Phase 1 configuration mainly defines the ends of the IPsec tunnel. The remote end is the remote gateway with which the FortiGate unit exchanges IPsec packets. The local end is the FortiGate interface that sends and receives IPsec packets.

If you want to control how the IKE negotiation is processed when there is no traffic, as well as the length of time the FortiGate unit waits for negotiations to occur, you can use the `negotiation-timeout` and `auto-negotiate` commands in the CLI.

For more information, refer to [Phase 2 parameters on page 1675](#) and [Phase 2 parameters on page 1675](#).

Name	<p>Type a name for the Phase 1 definition. The maximum name length is 15 characters for an interface mode VPN, 35 characters for a policy-based VPN. If Remote Gateway is Dialup User, the maximum name length is further reduced depending on the number of dialup tunnels that can be established: by 2 for up to 9 tunnels, by 3 for up to 99 tunnels, 4 for up to 999 tunnels, and so on.</p> <p>For a tunnel mode VPN, the name normally reflects where the remote connection originates. For a route-based tunnel, the FortiGate unit also uses the name for the virtual IPsec interface that it creates automatically.</p>
-------------	---

Remote Gateway	<p>Select the category of the remote connection:</p> <p>Static IP Address — If the remote peer has a static IP address.</p> <p>Dialup User — If one or more FortiClient or FortiGate dialup clients with dynamic IP addresses will connect to the FortiGate unit.</p> <p>Dynamic DNS — If a remote peer that has a domain name and subscribes to a dynamic DNS service will connect to the FortiGate unit.</p>
IP Address	If you selected Static IP Address , enter the IP address of the remote peer.
Dynamic DNS	If you selected Dynamic DNS , enter the domain name of the remote peer.
Local Interface	<p>This option is available in NAT mode only. Select the name of the interface through which remote peers or dialup clients connect to the FortiGate unit.</p> <p>By default, the local VPN gateway IP address is the IP address of the interface that you selected.</p>
Mode	<p>Main mode — the Phase 1 parameters are exchanged in multiple rounds with encrypted authentication information.</p> <p>Aggressive mode — the Phase 1 parameters are exchanged in single message with authentication information that is not encrypted.</p> <p>When the remote VPN peer has a dynamic IP address and is authenticated by a pre-shared key, you must select Aggressive mode if there is more than one dialup phase1 configuration for the interface IP address.</p> <p>When the remote VPN peer has a dynamic IP address and is authenticated by a certificate, you must select Aggressive mode if there is more than one Phase 1 configuration for the interface IP address and these Phase 1 configurations use different proposals.</p>
Authentication Method	Select Preshared Key or RSA Signature .
Pre-shared Key	<p>If you selected Pre-shared Key, enter the pre-shared key that the FortiGate unit will use to authenticate itself to the remote peer or dialup client during Phase 1 negotiations. You must define the same key at the remote peer or client.</p> <p>The key must contain at least 6 printable characters. For optimum protection against currently known attacks, the key must consist of a minimum of 16 randomly chosen alphanumeric characters. The limit is 128 characters.</p>

Certificate Name	If you selected RSA Signature , select the name of the server certificate that the FortiGate unit will use to authenticate itself to the remote peer or dialup client during Phase 1 negotiations. For information about obtaining and loading the required server certificate, see the FortiOS User Authentication guide .
Peer Options	Peer options are available to authenticate VPN peers or clients, depending on the Remote Gateway and Authentication Method settings.
Any peer ID	<p>Accept the local ID of any remote VPN peer or client. The FortiGate unit does not check identifiers (local IDs). You can set Mode to Aggressive or Main.</p> <p>You can use this option with RSA Signature authentication. But, for highest security, configure a PKI user/group for the peer and set Peer Options to Accept this peer certificate only.</p>
This peer ID	<p>This option is available when Aggressive Mode is enabled. Enter the identifier that is used to authenticate the remote peer. This identifier must match the Local ID that the remote peer's administrator has configured.</p> <p>If the remote peer is a FortiGate unit, the identifier is specified in the Local ID field of the Advanced Phase 1 configuration.</p> <p>If the remote peer is a FortiClient user, the identifier is specified in the Local ID field, accessed by selecting Config in the Policy section of the VPN connection's Advanced Settings.</p> <p>In circumstances where multiple remote dialup VPN tunnels exist, each tunnel must have a peer ID set.</p>
Peer ID from dialup group	<p>Authenticate multiple FortiGate or FortiClient dialup clients that use unique identifiers and unique pre-shared keys (or unique pre-shared keys only) through the same VPN tunnel.</p> <p>You must create a dialup user group for authentication purposes. Select the group from the list next to the Peer ID from dialup group option.</p> <p>You must set Mode to Aggressive when the dialup clients use unique identifiers and unique pre-shared keys. If the dialup clients use unique pre-shared keys only, you can set Mode to Main if there is only one dialup Phase 1 configuration for this interface IP address.</p>

Phase 1 advanced configuration settings

You can use the following advanced parameters to select the encryption and authentication algorithms that the FortiGate unit uses to generate keys for the IKE exchange. You can also use the following advanced parameters to ensure the smooth operation of Phase 1 negotiations.

These settings are mainly configured in the CLI, although some options are available after the tunnel is created using the VPN Creation Wizard (using the **Convert to Custom Tunnel** option).



If the FortiGate unit will act as a VPN client, and you are using security certificates for authentication, set the **Local ID** to the distinguished name (DN) of the local server certificate that the FortiGate unit will use for authentication purposes.

Note that, since FortiOS 5.4, an exact match is required to optimize IKE's gateway search utilizing binary trees. However, it is also possible to have partial matching of 'user.peer:cn' to match peers to gateways by performing a secondary match. When IKE receives IDi of type ASN1.DN, the first search is done with the whole DN string. If none is found, IKE will extract just the CN attribute value and perform a second search.

VXLAN over IPsec

Packets with VXLAN header are encapsulated within IPsec tunnel mode.

To configure VXLAN over IPsec - CLI:

```
config vpn ipsec phase1-interface/phase1
edit ipsec
set interface <name>
set encapsulation vxlan/gre
set encapsulation-address ike/ipv4/ipv6
set encap-local-gw4 xxx.xxx.xxx.xxx
set encap-remote-gw xxx.xxx.xxx.xxx
next
end
```

You can define an idle timer for IPsec tunnels. When no traffic has passed through the tunnel for the configured idle-timeout value, the IPsec tunnel will be flushed.

To configure IPsec tunnel idle timeout - CLI:

IPsec tunnel idle timer

```
config vpn ipsec phase1-interface
edit p1
set idle-timeout [enable | disable]
set idle-timeoutinterval <integer> //IPsec tunnel
idle timeout in minutes (10 - 43200).
end
end
```

IPv6 Version

Select if you want to use IPv6 addresses for the remote gateway and interface IP addresses.

Specify an IP address for the local end of the VPN tunnel. Select one of the following:

Local Gateway IP

Main Interface IP — The FortiGate unit obtains the IP address of the interface from the network interface settings.

Specify — Enter a secondary address of the interface selected in the Phase 1 **Local Interface** field.

You cannot configure Interface mode in a transparent mode VDOM.

Phase 1 Proposal

Select the encryption and authentication algorithms used to generate keys for protecting negotiations and add encryption and authentication algorithms as required.

You need to select a minimum of one and a maximum of three combinations. The remote peer or client must be configured to use at least one of the proposals that you define.

Select one of the following symmetric-key encryption algorithms:

DES — Digital Encryption Standard, a 64-bit block algorithm that uses a 56-bit key.

3DES — Triple-DES; plain text is encrypted three times by three keys.

AES128 — A 128-bit block algorithm that uses a 128-bit key.

AES192 — A 128-bit block algorithm that uses a 192-bit key.

AES256 — A 128-bit block algorithm that uses a 256-bit key.

ChaCha20/Poly1305 — A 128-bit block algorithm that uses a 128-bit key and a symmetric cipher. Only available for IKEv2.

You can select either of the following message digests to check the authenticity of messages during an encrypted session:

MD5 — Message Digest 5.

SHA1 — Secure Hash Algorithm 1 - a 160-bit message digest.

To specify one combination only, set the **Encryption** and **Authentication** options of the second combination to NULL. To specify a third combination, use the **Add** button beside the fields for the second combination.

Diffie-Hellman Group

Select one or more Diffie-Hellman groups from DH groups 1, 2, 5, and 14 through 21. At least one of the **Diffie-Hellman Group** settings on the remote peer or client must match one the selections on the FortiGate unit. Failure to match one or more DH groups will result in failed negotiations.

Keylife

Enter the time (in seconds) that must pass before the IKE encryption key expires. When the key expires, a new key is generated without interrupting service. The keylife can be from 120 to 172 800 seconds.

Local ID

If the FortiGate unit will act as a VPN client and you are using peer IDs for authentication purposes, enter the identifier that the FortiGate unit will supply to the VPN server during the Phase 1 exchange.

If the FortiGate unit will act as a VPN client, and you are using security certificates for authentication, select the distinguished name (DN) of the local server certificate that the FortiGate unit will use for authentication purposes.

If the FortiGate unit is a dialup client and will not be sharing a tunnel with other dialup clients (that is, the tunnel will be dedicated to this Fortinet dialup client), set **Mode** to **Aggressive**.

Note that this Local ID value must match the peer ID value given for the remote VPN peer's Peer Options.

XAuth

This option supports the authentication of dialup clients. It is available for IKE v1 only.

Disable — Select if you do not use XAuth.

Enable as Client — If the FortiGate unit is a dialup client, enter the user name and password that the FortiGate unit will need to authenticate itself to the remote XAuth server.

Enable as Server — This is available only if **Remote Gateway** is set to **Dialup User**. Dialup clients authenticate as members of a dialup user group. You must first create a user group for the dialup clients that need access to the network behind the FortiGate unit.

You must also configure the FortiGate unit to forward authentication requests to an external RADIUS or LDAP authentication server.

Select a **Server Type** setting to determine the type of encryption method to use between the FortiGate unit, the XAuth client and the external authentication server, and then select the user group from the User Group list.

Username

Enter the user name that is used for authentication.

Password

Enter the password that is used for authentication.

NAT Traversal

Select the check box if a NAT device exists between the local FortiGate unit and the VPN peer or client. The local FortiGate unit and the VPN peer or client must have the same NAT traversal setting (both selected or both cleared) to connect reliably.

Additionally, you can force IPsec to use NAT traversal. If NAT is set to **Forced**, the FortiGate will use a port value of zero when constructing the NAT discovery hash for the peer. This causes the peer to think it is behind a NAT device, and it will use UDP encapsulation for IPsec, even if no NAT is present. This approach maintains interoperability with any IPsec implementation that supports the NAT-T RFC.

Keepalive Frequency

If you enabled **NAT-traversal**, enter a keepalive frequency setting.

Dead Peer Detection

Select this check box to reestablish VPN tunnels on idle connections and clean up dead IKE peers if required. You can use this option to receive notification whenever a tunnel goes up or down, or to keep the tunnel connection open when no traffic is being generated inside the tunnel. For example, in scenarios where a dialup client or dynamic DNS peer connects from an IP address that changes periodically, traffic may be suspended while the IP address changes.

With **Dead Peer Detection** selected, you can use the `config vpn ipsec phase1 (tunnel mode)` or `config vpn ipsec phase1-interface (interface mode)` CLI command to optionally specify a retry count and a retry interval.

IKEv1 fragmentation

UDP fragmentation can cause issues in IPsec when either the ISP or perimeter firewall(s) cannot pass or fragment the oversized UDP packets that occur when using a very large public security key (PSK). The result is that IPsec tunnels do not come up. The solution is IKE fragmentation.

For most configurations, enabling IKE fragmentation allows connections to automatically establish when they otherwise might have failed due to intermediate nodes dropping IKE messages containing large certificates, which typically push the packet size over 1500 bytes.

FortiOS will fragment a packet on sending if, and only if, all the following are true:

- Phase 1 contains "set fragmentation enable".
- The packet is larger than the minimum MTU (576 for IPv4, 1280 for IPv6).
- The packet is being re-transmitted.

By default, IKE fragmentation is enabled, but upon upgrading, any existing phase1-interface may have have "set fragmentation disable" added in order to preserve the existing behaviour of not supporting fragmentation.

Enabling or disabling IKE fragmentation - CLI

```
config vpn ipsec phase1-interface
  edit 1
    set fragmentation [enable | disable]
  next
end
```

IKEv2 fragmentation

With IKEv2, because [RFC 7383](#) requires each fragment to be individually encrypted and authenticated, we would have to keep a copy of the unencrypted payloads around for each outgoing packet, in case the original single packet was never answered and we wanted to retry with fragments. With the following implementation, if the IKE payloads are greater than a configured threshold, the IKE packets are preemptively fragmented and encrypted.

CLI syntax

```
config vpn ipsec phase1-interface
    edit ike
        set ike-version 2
        set fragmentation [enable|disable]
        set fragmentation-mtu [500-16000]
    next
end
```

Phase 2 configuration

After IPsec Phase 1 negotiations end successfully, you begin Phase 2. You can configure the Phase 2 parameters to define the algorithms that the FortiGate unit may use to encrypt and transfer data for the remainder of the session. During Phase 2, you select specific IPsec security associations needed to implement security services and establish a tunnel.

The basic Phase 2 settings associate IPsec Phase 2 parameters with the Phase 1 configuration that specifies the remote end point of the VPN tunnel. In most cases, you need to configure only basic Phase 2 settings.

These settings are mainly configured in the CLI, although some options are available after the tunnel is created using the VPN Creation Wizard (using the **Convert to Custom Tunnel** option).

Name	Type a name to identify the Phase 2 configuration.
Phase 1	Select the Phase 1 tunnel configuration. For more information on configuring Phase 1, see Phase 1 configuration on page 1642 . The Phase 1 configuration describes how remote VPN peers or clients will be authenticated on this tunnel, and how the connection to the remote peer or client will be secured.
Advanced	Define advanced Phase 2 parameters. For more information, see Phase 2 advanced configuration settings below.

Phase 2 advanced configuration settings

In Phase 2, the FortiGate unit and the VPN peer or client exchange keys again to establish a secure communication channel between them. You select the encryption and authentication algorithms needed to generate keys for protecting the implementation details of Security Associations (SAs). These are called Phase 2 Proposal parameters. The keys are generated automatically using a Diffie-Hellman algorithm.

You can use a number of additional advanced Phase 2 settings to enhance the operation of the tunnel.

Phase 2 Proposal	<p>Select the encryption and authentication algorithms that will be proposed to the remote VPN peer. You can specify up to three proposals. To establish a VPN connection, at least one of the proposals that you specify must match configuration on the remote peer.</p> <p>Initially there are two proposals. Add and Delete icons are next to the second Authentication field.</p> <p>It is invalid to set both Encryption and Authentication to NULL.</p>
Encryption	<p>Select a symmetric-key algorithms:</p> <p>NULL — Do not use an encryption algorithm.</p> <p>DES — Digital Encryption Standard, a 64-bit block algorithm that uses a 56-bit key.</p> <p>3DES — Triple-DES; plain text is encrypted three times by three keys.</p> <p>AES128 — A 128-bit block algorithm that uses a 128-bit key.</p> <p>AES192 — A 128-bit block algorithm that uses a 192-bit key.</p> <p>AES256 — A 128-bit block algorithm that uses a 256-bit key.</p> <p>ChaCha20/Poly1305 — A 128-bit block algorithm that uses a 128-bit key and a symmetric cipher. Only available for IKEv2.</p>
Authentication	<p>You can select either of the following message digests to check the authenticity of messages during an encrypted session:</p> <p>NULL — Do not use a message digest.</p> <p>MD5 — Message Digest 5.</p> <p>SHA1 — Secure Hash Algorithm 1 - a 160-bit message digest.</p> <p>To specify one combination only, set the Encryption and Authentication options of the second combination to NULL. To specify a third combination, use the Add button beside the fields for the second combination.</p>
Enable replay detection	Replay attacks occur when an unauthorized party intercepts a series of IPsec packets and replays them back into the tunnel.
Enable perfect forward secrecy (PFS)	Perfect forward secrecy (PFS) improves security by forcing a new Diffie-Hellman exchange whenever keylife expires.
Diffie-Hellman Group	Select one Diffie-Hellman group (1, 2, 5, or 14 through 21). This must match the DH Group that the remote peer or dialup client uses.
Keylife	<p>Select the method for determining when the Phase 2 key expires: Seconds, KBytes, or Both. If you select Both, the key expires when either the time has passed or the number of KB have been processed.</p>
Autokey Keep Alive	Select the check box if you want the tunnel to remain active when no data is being processed.

Auto-negotiate	Enable the option if you want the tunnel to be automatically renegotiated when the tunnel expires.
DHCP-IPsec	<p>Provide IP addresses dynamically to VPN clients. This is available for Phase 2 configurations associated with a dialup Phase 1 configuration.</p> <p>You also need configure a DHCP server or relay on the private network interface. You must configure the DHCP parameters separately.</p> <p>If you configure the DHCP server to assign IP addresses based on RADIUS user group attributes, you must also set the Phase 1 Peer Options to Peer ID from dialup group and select the appropriate user group. See Phase 1 configuration on page 1642.</p> <p>If the FortiGate unit acts as a dialup server and you manually assigned FortiClient dialup clients VIP addresses that match the network behind the dialup server, selecting the check box will cause the FortiGate unit to act as a proxy for the dialup clients.</p>
Quick Mode Selector	<p>Specify the source and destination IP addresses to be used as selectors for IKE negotiations. If the FortiGate unit is a dialup server, keep the default value of 0.0.0.0/0 unless you need to circumvent problems caused by ambiguous IP addresses between one or more of the private networks making up the VPN. You can specify a single host IP address, an IP address range, or a network address. You may optionally specify source and destination port numbers and a protocol number.</p> <p>If you are editing an existing Phase 2 configuration, the Source address and Destination address fields are unavailable if the tunnel has been configured to use firewall addresses as selectors. This option exists only in the CLI.</p>
Source address	<p>If the FortiGate unit is a dialup server, enter the source IP address that corresponds to the local senders or network behind the local VPN peer (for example, 172.16.5.0/24 or 172.16.5.0/255.255.255.0 for a subnet, or 172.16.5.1/32 or 172.16.5.1/255.255.255.255 for a server or host, or 192.168.10.[80-100] or 192.168.10.80-192.168.10.100 for an address range). A value of 0.0.0.0/0 means all IP addresses behind the local VPN peer.</p> <p>If the FortiGate unit is a dialup client, source address must refer to the private network behind the Fortinet dialup client.</p>
Source port	Enter the port number that the local VPN peer uses to transport traffic related to the specified service (protocol number). The range is from 0 to 65535. To specify all ports, type 0.

Destination address	Enter the destination IP address that corresponds to the recipients or network behind the remote VPN peer (for example, 192.168.20.0/24 for a subnet, or 172.16.5.1/32 for a server or host, or 192.168.10.[80-100] for an address range). A value of 0.0.0.0/0 means all IP addresses behind the remote VPN peer.
Destination port	Enter the port number that the remote VPN peer uses to transport traffic related to the specified service (protocol number). To specify all ports, enter 0.
Protocol	Enter the IP protocol number of the service. To specify all services, enter 0.

FortiClient VPN

Use the **FortiClient VPN for OS X, Windows, and Android** VPN Wizard option when configuring an IPsec VPN for remote users to connect to the VPN tunnel using FortiClient.

When configuring a FortiClient VPN connection, the settings for Phase 1 and Phase 2 settings are automatically configured by the FortiGate unit. They are set to:

- Remote Gateway — Dialup User
- Mode — Aggressive
- Default settings for Phase 1 and 2 Proposals
- XAUTH Enable as Server (Auto)
- IKE mode-config will be enabled
- Peer Option — “Any peer ID”

The remainder of the settings use the current FortiGate defaults. Note that FortiClient settings need to match these FortiGate defaults. If you need to configure advanced settings for the FortiClient VPN, you must do so using the CLI.

Name	Enter a name for the FortiClient VPN.
Local Outgoing Interface	Select the local outgoing interface for the VPN.
Authentication Method	Select the type of authentication used when logging in to the VPN.
Preshared Key	If Pre-shared Key was selected in Authentication Method , enter the pre-shared key in the field provided.
User Group	Select a user group. You can also create a user group from the drop-down list by selecting Create New .
Address Range Start IP	Enter the start IP address for the DHCP address range for the client.
Address Range End IP	Enter the end IP address for the address range.
Subnet Mask	Enter the subnet mask.

Enable IPv4 Split Tunnel	Enabled by default, this option enables the FortiClient user to use the VPN to access internal resources while other Internet access is not sent over the VPN, alleviating potential traffic bottlenecks in the VPN connection. Disable this option to have all traffic sent through the VPN tunnel.
Accessible Networks	Select from a list of internal networks that the FortiClient user can access.
Client Options	<p>These options affect how the FortiClient application behaves when connected to the FortiGate VPN tunnel. When enabled, a check box for the corresponding option appears on the VPN login screen in FortiClient, and is not enabled by default.</p> <p>Save Password - When enabled, if the user selects this option, their password is stored on the user's computer and will automatically populate each time they connect to the VPN.</p> <p>Auto Connect - When enabled, if the user selects this option, when the FortiClient application is launched, for example after a reboot or system startup, FortiClient will automatically attempt to connect to the VPN tunnel.</p> <p>Always Up (Keep Alive) - When enabled, if the user selects this option, the FortiClient connection will not shut down. When not selected, during periods of inactivity, FortiClient will attempt to stay connected every three minutes for a maximum of 10 minutes.</p>
Endpoint Registration	<p>When selected, the FortiGate unit requests a registration key from FortiClient before a connection can be established. A registration key is defined by going to System > Advanced.</p> <p>For more information on FortiClient VPN connections to a FortiGate unit, see the FortiClient Administration Guide.</p>
DNS Server	<p>Select which DNS server to use for this VPN:</p> <p>Use System DNS — Use the same DNS servers as the FortiGate unit. These are configured at Network > DNS. This is the default option.</p> <p>Specify — Specify the IP address of a different DNS server.</p>

Concentrator

In a hub-and-spoke configuration, policy-based VPN connections to a number of remote peers radiate from a single, central FortiGate unit. Site-to-site connections between the remote peers do not exist; however, you can establish VPN tunnels between any two of the remote peers through the FortiGate unit's "hub".

In a hub-and-spoke network, all VPN tunnels terminate at the hub. The peers that connect to the hub are known as "spokes". The hub functions as a concentrator on the network, managing all VPN connections between the spokes. VPN traffic passes from one tunnel to the other through the hub.

You define a concentrator to include spokes in the hub-and-spoke configuration. You create the concentrator in **VPN > IPsec Concentrator** and select **Create New**. A concentrator configuration specifies which spokes to include in an IPsec hub-and-spoke configuration.

Concentrator Name	Type a name for the concentrator.
Available Tunnels	A list of defined IPsec VPN tunnels. Select a tunnel from the list and then select the right arrow.
Members	A list of tunnels that are members of the concentrator. To remove a tunnel from the concentrator, select the tunnel and select the left arrow.

IPsec Monitor

You can use the IPsec Monitor to view activity on IPsec VPN tunnels and start or stop those tunnels. The display provides a list of addresses, proxy IDs, and timeout information for all active tunnels, including tunnel mode and route-based (interface mode) tunnels.

To view the IPsec monitor, go to **Monitor > IPsec Monitor**.



Tunnels are considered as "up" if at least one phase 2 selector is active. To avoid confusion, when a tunnel is down, **IPsec Monitor** will keep the **Phase 2 Selectors** column, but hide it by default and be replaced with **Phase 1** status column.

For dialup VPNs, the list provides status information about the VPN tunnels established by dialup clients, and their IP addresses.

For static IP or dynamic DNS VPNs, the list provides status and IP addressing information about VPN tunnels, active or not, to remote peers that have static IP addresses or domain names. You can also start and stop individual tunnels from the list.

Timeout field in IPsec Monitor page

The **Timeout** field in **Monitor > IPsec Monitor** shows the realtime timeout value for each VPN tunnel that is **Up** (tunnels that are **Down** show a timeout value of **0**).

Phase 1 parameters

This chapter provides detailed step-by-step procedures for configuring a FortiGate unit to accept a connection from a remote peer or dialup client. The Phase 1 parameters identify the remote peer or clients and supports authentication through preshared keys or digital certificates. You can increase access security further using peer identifiers, certificate distinguished names, group names, or the FortiGate extended authentication (XAuth) option for authentication purposes.

For more information on Phase 1 parameters in the web-based manager, see [IPsec VPN in the web-based manager on page 1642](#).

The information and procedures in this section do not apply to VPN peers that perform negotiations using manual keys.

The following topics are included in this section:

Overview

To configure IPsec Phase 1 settings, go to **VPN > IPsec Tunnels** and edit the **Phase 1 Proposal** (if it is not available, you may need to click the **Convert to Custom Tunnel** button).

IPsec Phase 1 settings define:

- The remote and local ends of the IPsec tunnel
- If Phase 1 parameters are exchanged in multiple rounds with encrypted authentication information (main mode) or in a single message with authentication information that is not encrypted (aggressive mode)
- If a preshared key or digital certificates will be used to authenticate the FortiGate unit to the VPN peer or dialup client
- If the VPN peer or dialup client is required to authenticate to the FortiGate unit. A remote peer or dialup client can authenticate by peer ID or, if the FortiGate unit authenticates by certificate, it can authenticate by peer certificate.
- The IKE negotiation proposals for encryption and authentication
- Optional XAuth authentication, which requires the remote user to enter a user name and password. A FortiGate VPN server can act as an XAuth server to authenticate dialup users. A FortiGate unit that is a dialup client can also be configured as an XAuth client to authenticate itself to the VPN server.

For all the Phase 1 web-based manager fields, see [IPsec VPN in the web-based manager on page 1642](#).

If you want to control how IKE is negotiated when there is no traffic, as well as the length of time the unit waits for negotiations to occur, use the `negotiation-timeout` and `auto-negotiate` commands in the CLI.

Defining the tunnel ends

To begin defining the Phase 1 configuration, go to **VPN > IPsec Tunnels** and select **Create New**. Enter a unique descriptive name for the VPN tunnel and follow the instructions in the VPN Creation Wizard.

The Phase 1 configuration mainly defines the ends of the IPsec tunnel. The remote end is the remote gateway with which the FortiGate unit exchanges IPsec packets. The local end is the FortiGate interface that sends and receives IPsec packets.

The remote gateway can be:

- A static IP address
- A domain name with a dynamic IP address
- A dialup client

A statically addressed remote gateway is the simplest to configure. You specify the IP address. Unless restricted in the security policy, either the remote peer or a peer on the network behind the FortiGate unit can bring up the tunnel.

If the remote peer has a domain name and subscribes to a dynamic DNS service, you need to specify only the domain name. The FortiGate unit performs a DNS query to determine the appropriate IP address. Unless restricted in the security policy, either the remote peer or a peer on the network behind the FortiGate unit can bring up the tunnel.

If the remote peer is a dialup client, only the dialup client can bring up the tunnel. The IP address of the client is not known until it connects to the FortiGate unit. This configuration is a typical way to provide a VPN for client PCs running VPN client software such as the FortiClient Endpoint Security application.

The local end of the VPN tunnel, the Local Interface, is the FortiGate interface that sends and receives the IPsec packets. This is usually the public interface of the FortiGate unit that is connected to the Internet (typically the WAN1 port). Packets from this interface pass to the private network through a security policy.

By default, the local VPN gateway is the IP address of the selected Local Interface. If you are configuring an interface mode VPN, you can optionally use a secondary IP address of the Local Interface as the local gateway.

Choosing Main mode or Aggressive mode

The FortiGate unit and the remote peer or dialup client exchange Phase 1 parameters in either Main mode or Aggressive mode. This choice does not apply if you use IKE version 2, which is available only for route-based configurations.

- In **Main** mode, the Phase 1 parameters are exchanged in multiple rounds with encrypted authentication information
- In **Aggressive** mode, the Phase 1 parameters are exchanged in a single message with unencrypted authentication information.

Although Main mode is more secure, you must select Aggressive mode if there is more than one dialup Phase 1 configuration for the interface IP address, and the remote VPN peer or client is authenticated using an identifier local ID. Aggressive mode might not be as secure as Main mode, but the advantage to Aggressive mode is that it is faster than Main mode (since fewer packets are exchanged). Aggressive mode is typically used for remote access VPNs. But you would also use aggressive mode if one or both peers have dynamic external IP addresses. Descriptions of the peer options in this guide indicate whether Main or Aggressive mode is required.

Choosing the IKE version

If you create a route-based VPN, you have the option of selecting IKE version 2. Otherwise, IKE version 1 is used.

IKEv2, defined in [RFC 4306](#), simplifies the negotiation process that creates the security association (SA).

If you select IKEv2:

- There is no choice in Phase 1 of Aggressive or Main mode.
- Extended Authentication (XAUTH) is not available.
- You can select only one Diffie-Hellman Group.
- You can utilize EAP and MOBIKE.

Repeated authentication in IKEv2

This feature provides the option to control whether a device requires its peer to re-authenticate or whether re-key is sufficient. It does not influence the re-authentication or re-key behavior of the device itself, which is controlled by the peer (with the default being to re-key). This solution is in response to [RFC 4478](#). As described by the IETF, "the purpose of this is to limit the time that security associations (SAs) can be used by a third party who has gained control of the IPsec peer".

Syntax

```
config vpn ipsec phase1-interface
    edit p1
        set reauth [enable | disable]
    next
end
```

IKEv2 cookie notification for IKE_SA_INIT

IKEv2 offers an optional exchange within IKE_SA_INIT (the initial exchange between peers when establishing a secure tunnel) as a result of an inherent vulnerability in IPsec implementations, as described in [RFC 5996](#).

Two expected attacks against IKE are state and CPU exhaustion, where the target is flooded with session initiation requests from forged IP addresses. These attacks can be made less effective if a responder uses minimal CPU and commits no state to an SA until it knows the initiator can receive packets at the address from which it claims to be sending them.

If the IKE_SA_INIT response includes the cookie notification, the initiator MUST then retry the IKE_SA_INIT request, and include the cookie notification containing the received data as the first payload, and all other payloads unchanged.

Upon detecting that the number of half-open IKEv2 SAs is above the threshold value, the VPN dialup server requires all future SA_INIT requests to include a valid cookie notification payload that the server sends back, in order to preserve CPU and memory resources.

For most devices, the threshold value is set to 500, half of the maximum 1,000 connections.

This feature is enabled by default in FortiOS 5.4.

IKEv2 Quick Crash Detection

There is support for IKEv2 Quick Crash Detection (QCD) as described in [RFC 6290](#).

RFC 6290 describes a method in which an IKE peer can quickly detect that the gateway peer that it has and established an IKE session with has rebooted, crashed, or otherwise lost IKE state. When the gateway receives IKE messages or ESP packets with unknown IKE or IPsec SPIs, the IKEv2 protocol allows the gateway to send the peer an unprotected IKE message containing INVALID_IKE_SPI or INVALID_SPI notification payloads.

RFC 6290 introduces the concept of a QCD token, which is generated from the IKE SPIs and a private QCD secret, and exchanged between peers during the protected IKE AUTH exchange.

Adding Quick Crash Detection - CLI Syntax

```
config system settings
    set ike-quick-crash-detect [enable | disable]
end
```


IKEv1 Quick Crash Detection

Based on the IKEv2 QCD feature described above, IKEv1 QCD is implemented using a new IKE vendor ID, "Fortinet Quick Crash Detection", and so both endpoints must be FortiGate devices. The QCD token is sent in the Phase 1 exchange and must be encrypted, so this is only implemented for IKEv1 in Main mode (Aggressive mode is not supported as there is no available AUTH message in which to include the token).

Otherwise, the feature works the same as in IKEv2 (RFC 6290).

Authenticating the FortiGate unit

The FortiGate unit can authenticate itself to remote peers or dialup clients using either a pre-shared key or an RSA Signature (certificate).

Authenticating the FortiGate unit with digital certificates

To authenticate the FortiGate unit using digital certificates, you must have the required certificates installed on the remote peer and on the FortiGate unit. The signed server certificate on one peer is validated by the presence of the root certificate installed on the other peer. If you use certificates to authenticate the FortiGate unit, you can also require the remote peers or dialup clients to authenticate using certificates.

For more information about obtaining and installing certificates, see the [FortiOS User Authentication guide](#).

Authenticating the FortiGate unit using digital certificates

1. Go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
2. Edit the **Phase 1 Proposal** (if it is not available, you may need to click the **Convert to Custom Tunnel** button):

Name	Enter a name that reflects the origination of the remote connection. For interface mode, the name can be up to 15 characters long.
Remote Gateway	<p>Select the nature of the remote connection.</p> <p>Each option changes the available fields you must configure. For more information, see Authenticating the FortiGate unit on page 1658.</p>
Local Interface	Select the interface that is the local end of the IPsec tunnel. For more information, see Authenticating the FortiGate unit on page 1658 . The local interface is typically the WAN1 port.
Mode	<p>Select a mode. It is easier to use Aggressive mode.</p> <p>In Main mode, parameters are exchanged in multiple encrypted rounds.</p> <p>In Aggressive mode, parameters are exchanged in a single unencrypted message.</p> <p>Aggressive mode must be used when the remote VPN peer or client has a dynamic IP address, or the remote VPN peer or client will be authenticated using an identifier (local ID).</p> <p>For more information, see Authenticating the FortiGate unit on page 1658.</p>

Authentication Method	Select Signature .
Certificate Name	<p>Select the name of the server certificate that the FortiGate unit will use to authenticate itself to the remote peer or dialup client during Phase 1 negotiations.</p> <p>You must obtain and load the required server certificate before this selection. See the FortiOS User Authentication guide. If you have not loaded any certificates, use the certificate named Fortinet_Factory.</p>
Peer Options	<p>Peer options define the authentication requirements for remote peers or dialup clients. They are not for your FortiGate unit itself.</p> <p>See Authenticating the FortiGate unit on page 1658.</p>
Advanced	<p>You can use the default settings for most Phase 1 configurations. Changes are required only if your network requires them. These settings includes IKE version, DNS server, P1 proposal encryption and authentication settings, and XAuth settings. See Authenticating the FortiGate unit on page 1658.</p>

3. If you are configuring authentication parameters for a dialup user group, optionally define extended authentication (XAuth) parameters in the Advanced section. See [Authenticating the FortiGate unit on page 1658](#).
4. Select **OK**.

Authenticating the FortiGate unit with a pre-shared key

The simplest way to authenticate a FortiGate unit to its remote peers or dialup clients is by means of a pre-shared key. This is less secure than using certificates, especially if it is used alone, without requiring peer IDs or extended authentication (XAuth). Also, you need to have a secure way to distribute the pre-shared key to the peers.

If you use pre-shared key authentication alone, all remote peers and dialup clients must be configured with the same pre-shared key. Optionally, you can configure remote peers and dialup clients with unique pre-shared keys. On the FortiGate unit, these are configured in user accounts, not in the phase_1 settings. For more information, see [Authenticating the FortiGate unit on page 1658](#).

The pre-shared key must contain at least 6 printable characters and best practices dictate that it be known only to network administrators. For optimum protection against currently known attacks, the key must consist of a minimum of 16 randomly chosen alphanumeric characters.

If you authenticate the FortiGate unit using a pre-shared key, you can require remote peers or dialup clients to authenticate using peer IDs, but not client certificates.

Authenticating the FortiGate unit with a pre-shared key

1. Go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
2. Edit the **Phase 1 Proposal** (if it is not available, you may need to click the **Convert to Custom Tunnel** button):

Name	Enter a name that reflects the origination of the remote connection.
-------------	--

Remote Gateway	Select the nature of the remote connection. For more information, see Authenticating the FortiGate unit on page 1658 .
Local Interface	Select the interface that is the local end of the IPsec tunnel. For more information, see Authenticating the FortiGate unit on page 1658 . The local interface is typically the WAN1 port.
Mode	<p>Select Main or Aggressive mode.</p> <p>In Main mode, the Phase 1 parameters are exchanged in multiple rounds with encrypted authentication information.</p> <p>In Aggressive mode, the Phase 1 parameters are exchanged in single message with authentication information that is not encrypted.</p> <p>When the remote VPN peer or client has a dynamic IP address, or the remote VPN peer or client will be authenticated using an identifier (local ID), you must select Aggressive mode if there is more than one dialup Phase 1 configuration for the interface IP address.</p> <p>For more information, see Authenticating the FortiGate unit on page 1658.</p>
Authentication Method	Select Pre-shared Key .
Pre-shared Key	Enter the preshared key that the FortiGate unit will use to authenticate itself to the remote peer or dialup client during Phase 1 negotiations. You must define the same value at the remote peer or client. The key must contain at least 6 printable characters and best practices dictate that it only be known by network administrators. For optimum protection against currently known attacks, the key must consist of a minimum of 16 randomly chosen alphanumeric characters.
Peer options	Peer options define the authentication requirements for remote peers or dialup clients, not for the FortiGate unit itself. You can require the use of peer IDs, but not client certificates. For more information, see Authenticating the FortiGate unit on page 1658 .
Advanced	You can retain the default settings unless changes are needed to meet your specific requirements. See Authenticating the FortiGate unit on page 1658 .

3. If you are configuring authentication parameters for a dialup user group, optionally define extended authentication (XAuth) parameters. See [Authenticating the FortiGate unit on page 1658](#).
4. Select **OK**.

Authenticating remote peers and clients

Certificates or pre-shared keys restrict who can access the VPN tunnel, but they do not identify or authenticate the remote peers or dialup clients. You have the following options for authentication:

Methods of authenticating remote VPN peers

Certificates or Pre-shared key	Local ID	User account pre-shared keys	Reference
Certificates			See Enabling VPN access for specific certificate holders on page 1661 .
Either	X		See Enabling VPN access by peer identifier on page 1663 .
Pre-shared key		X	See Enabling VPN access with user accounts and pre-shared keys on page 1664 .
Pre-shared key	X	X	See Enabling VPN access with user accounts and pre-shared keys on page 1664 .

Repeated authentication in Internet Key Exchange (IKEv2) protocol

This feature provides the option to control whether a device requires its peer to re-authenticate or whether re-key is sufficient. It does not influence the re-authentication or re-key behavior of the device itself, which is controlled by the peer (with the default being to re-key).

This solution is in response to [RFC 4478](#). This solution is intended to limit the time that security associations (SAs) can be used by a third party who has gained control of the IPsec peer.

CLI syntax:

```
config vpn ipsec phase1-interface
edit p1
    set reauth [enable | disable]
next
end
```

disable: Disable IKE SA re-authentication.

enable: Enable IKE SA re-authentication.

Enabling VPN access for specific certificate holders

When a VPN peer or dialup client is configured to authenticate using digital certificates, it sends the Distinguished Name (DN) of its certificate to the FortiGate unit. This DN can be used to allow VPN access for the certificate holder. That is, a FortiGate unit can be configured to deny connections to all remote peers and dialup clients except the one having the specified DN.

Before you begin

The following procedures assume that you already have an existing Phase 1 configuration (see [Authenticating remote peers and clients on page 1660](#)). Follow the procedures below to add certificate-based authentication parameters to the existing configuration.

Before you begin, you must obtain the certificate DN of the remote peer or dialup client. If you are using the FortiClient application as a dialup client, refer to FortiClient online help for information about how to view the certificate DN. To view the certificate DN of a FortiGate unit, see [Viewing server certificate information and obtaining the local DN on page 1662](#).

Use the `config user peer` CLI command to load the DN value into the FortiGate configuration. For example, if a remote VPN peer uses server certificates issued by your own organization, you would enter information similar to the following:

```
config user peer
  edit DN_FG1000
    set cn 192.168.2.160
    set cn-type ipv4
  end
```

The value that you specify to identify the entry (for example, DN_FG1000) is displayed in the **Accept this peer certificate only** list in the IPsec Phase 1 configuration when you return to the web-based manager.

If the remote VPN peer has a CA-issued certificate to support a higher level of credibility, you would enter information similar to the following in the CLI:

```
config user peer
  edit CA_FG1000
    set ca CA_Cert_1
    set subject FG1000_at_site1
  end
```

The value that you specify to identify the entry (for example, CA_FG1000) is displayed in the **Accept this peer certificate only** list in the IPsec Phase 1 configuration when you return to the web-based manager. For more information about these CLI commands, see the “user” chapter of the [FortiGate CLI Reference](#).

A group of certificate holders can be created based on existing user accounts for dialup clients. To create the user accounts for dialup clients, see the “User” chapter of the [FortiGate Administration Guide](#). To create the certificate group afterward, use the `config user peergrp` CLI command. See the “user” chapter of the [FortiGate CLI Reference](#).

Viewing server certificate information and obtaining the local DN

1. Go to **System > Certificates**.
2. Note the CN value in the **Subject** field (for example, CN = 172.16.10.125, CN = info@fortinet.com, or CN = www.example.com).

Viewing CA root certificate information and obtaining the CA certificate name

1. Go to **System > Certificates > CA Certificates**.
2. Note the value in the **Name** column (for example, CA_Cert_1).

Configuring certificate authentication for a VPN

With peer certificates loaded, peer users and peer groups defined, you can configure your VPN to authenticate users by certificate.

Enabling access for a specific certificate holder or a group of certificate holders

1. At the FortiGate VPN server, go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
2. Edit the **Phase 1 Proposal** (if it is not available, you may need to click the **Convert to Custom Tunnel** button).
3. From the **Authentication Method** list, select **RSA Signature**.
4. From the **Certificate Name** list, select the name of the server certificate that the FortiGate unit will use to authenticate itself to the remote peer or dialup client
5. Under **Peer Options**, select one of these options:
 - To accept a specific certificate holder, select **Accept this peer certificate only** and select the name of the certificate that belongs to the remote peer or dialup client. The certificate DN must be added to the FortiGate configuration through CLI commands before it can be selected here. See [Before you begin on page 1662](#).
 - To accept dialup clients who are members of a certificate group, select **Accept this peer certificate group only** and select the name of the group. The group must be added to the FortiGate configuration through CLI commands before it can be selected here. See [Before you begin on page 1662](#).
6. If you want the FortiGate VPN server to supply the DN of a local server certificate for authentication purposes, select **Advanced** and then from the **Local ID** list, select the DN of the certificate that the FortiGate VPN server is to use.
7. Select **OK**.

Enabling VPN access by peer identifier

Whether you use certificates or pre-shared keys to authenticate the FortiGate unit, you can require that remote peers or clients have a particular peer ID. This adds another piece of information that is required to gain access to the VPN. More than one FortiGate/FortiClient dialup client may connect through the same VPN tunnel when the dialup clients share a preshared key and assume the same identifier.



In circumstances where multiple remote dialup VPN tunnels exist, each tunnel must have a peer ID set.

A peer ID, also called local ID, can be up to 63 characters long containing standard regular expression characters. Local ID is set in phase1 Aggressive Mode configuration.

You cannot require a peer ID for a remote peer or client that uses a pre-shared key and has a static IP address.

Authenticating remote peers or dialup clients using one peer ID

1. At the FortiGate VPN server, go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
2. Edit the **Phase 1 Proposal** (if it is not available, you may need to click the **Convert to Custom Tunnel** button).
3. Select **Aggressive** mode in any of the following cases:

- The FortiGate VPN server authenticates a FortiGate dialup client that uses a dedicated tunnel
 - A FortiGate unit has a dynamic IP address and subscribes to a dynamic DNS service
 - FortiGate/FortiClient dialup clients sharing the same preshared key and local ID connect through the same VPN tunnel
4. For the **Peer Options**, select **This peer ID** and type the identifier into the corresponding field.
 5. Select **OK**.

Assigning an identifier (local ID) to a FortiGate unit

Use this procedure to assign a peer ID to a FortiGate unit that acts as a remote peer or dialup client.

1. Go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
2. Edit the **Phase 1 Proposal** (if it is not available, you may need to click the **Convert to Custom Tunnel** button).
3. Select **Advanced**.
4. In the **Local ID** field, type the identifier that the FortiGate unit will use to identify itself.
5. Set **Mode** to **Aggressive** if any of the following conditions apply:
 - The FortiGate unit is a dialup client that will use a unique ID to connect to a FortiGate dialup server through a dedicated tunnel.
 - The FortiGate unit has a dynamic IP address, subscribes to a dynamic DNS service, and will use a unique ID to connect to the remote VPN peer through a dedicated tunnel.
 - The FortiGate unit is a dialup client that shares the specified ID with multiple dialup clients to connect to a FortiGate dialup server through the same tunnel.
6. Select **OK**.

Configuring the FortiClient application

Follow this procedure to add a peer ID to an existing FortiClient configuration:

1. Start the FortiClient application.
2. Go to **VPN > Connections**, select the existing configuration.
3. Select **Advanced > Edit > Advanced**.
4. Under **Policy**, select **Config**.
5. In the **Local ID** field, type the identifier that will be shared by all dialup clients. This value must match the **This peer ID** value that you specified previously in the Phase 1 gateway configuration on the FortiGate unit.
6. Select **OK** to close all dialog boxes.
7. Configure all dialup clients the same way using the same preshared key and local ID.

Enabling VPN access with user accounts and pre-shared keys

You can permit access only to remote peers or dialup clients that have pre-shared keys and/or peer IDs configured in user accounts on the FortiGate unit.

If you want two VPN peers (or a FortiGate unit and a dialup client) to accept reciprocal connections based on peer IDs, you must enable the exchange of their identifiers when you define the Phase 1 parameters.

The following procedures assume that you already have an existing Phase 1 configuration (see [Authenticating remote peers and clients on page 1660](#)). Follow the procedures below to add ID checking to the existing configuration.

Before you begin, you must obtain the identifier (local ID) of the remote peer or dialup client. If you are using the FortiClient Endpoint Security application as a dialup client, refer to the Authenticating FortiClient Dialup Clients Technical Note to view or assign an identifier. To assign an identifier to a FortiGate dialup client or a FortiGate unit that has a dynamic IP address and subscribes to a dynamic DNS service, see [Assigning an identifier \(local ID\) to a FortiGate unit on page 1664](#).

If required, a dialup user group can be created from existing user accounts for dialup clients. To create the user accounts and user groups, see the [User Authentication](#) handbook chapter.

The following procedure supports FortiGate/FortiClient dialup clients that use unique preshared keys and/or peer IDs. The client must have an account on the FortiGate unit and be a member of the dialup user group.

The dialup user group must be added to the FortiGate configuration before it can be selected. For more information, see the [User Authentication](#) handbook chapter.

The FortiGate dialup server compares the local ID that you specify at each dialup client to the FortiGate user-account user name. The dialup-client preshared key is compared to a FortiGate user-account password.

Authenticating dialup clients using unique preshared keys and/or peer IDs

1. At the FortiGate VPN server, go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
2. Edit the **Phase 1 Proposal** (if it is not available, you may need to click the **Convert to Custom Tunnel** button).
3. If the clients have unique peer IDs, set **Mode** to **Aggressive**.
4. Clear the **Pre-shared Key** field.
The user account password will be used as the preshared key.
5. Select **Peer ID from dialup group** and then select the group name from the list of user groups.
6. Select **OK**.

Follow this procedure to add a unique pre-shared key and unique peer ID to an existing FortiClient configuration.

Configuring FortiClient - pre-shared key and peer ID

1. Start the FortiClient Endpoint Security application.
2. Go to **VPN > Connections**, select the existing configuration.
3. Select **Advanced > Edit**.
4. In the **Preshared Key** field, type the FortiGate password that belongs to the dialup client (for example, 1234546).
The user account password will be used as the preshared key.
5. Select **Advanced**.
6. Under **Policy**, select **Config**.
7. In the **Local ID** field, type the FortiGate user name that you assigned previously to the dialup client (for example, FortiClient1).
8. Select **OK** to close all dialog boxes.

Configure all FortiClient dialup clients this way using unique preshared keys and local IDs.

Follow this procedure to add a unique pre-shared key to an existing FortiClient configuration.

Configuring FortiClient - preshared key only

1. Start the FortiClient Endpoint Security application.
2. Go to **VPN > Connections**, select the existing configuration

3. Select **Advanced > Edit**.
4. In the **Preshared Key** field, type the user name, followed by a “+” sign, followed by the password that you specified previously in the user account settings on the FortiGate unit (for example, FC2+1FG6LK).
5. Select **OK** to close all dialog boxes.

Configure all the FortiClient dialup clients this way using their unique peer ID and pre-shared key values.

Defining IKE negotiation parameters

In Phase 1, the two peers exchange keys to establish a secure communication channel between them. As part of the Phase 1 process, the two peers authenticate each other and negotiate a way to encrypt further communications for the duration of the session. The Phase 1 Proposal parameters select the encryption and authentication algorithms that are used to generate keys for protecting negotiations.

The IKE negotiation parameters determine:

- Which encryption algorithms may be applied for converting messages into a form that only the intended recipient can read
- Which authentication hash may be used for creating a keyed hash from a preshared or private key
- Which Diffie-Hellman group (DH Group) will be used to generate a secret session key

Phase 1 negotiations (in main mode or aggressive mode) begin as soon as a remote VPN peer or client attempts to establish a connection with the FortiGate unit. Initially, the remote peer or dialup client sends the FortiGate unit a list of potential cryptographic parameters along with a session ID. The FortiGate unit compares those parameters to its own list of advanced Phase 1 parameters and responds with its choice of matching parameters to use for authenticating and encrypting packets. The two peers handle the exchange of encryption keys between them, and authenticate the exchange through a preshared key or a digital signature.

Generating keys to authenticate an exchange

The FortiGate unit supports the generation of secret session keys automatically using a Diffie-Hellman algorithm. These algorithms are defined in [RFC 2409](#). The **Keylife** setting in the **Phase 1 Proposal** area determines the amount of time before the Phase 1 key expires. Phase 1 negotiations are re-keyed automatically when there is an active security association. See [Dead Peer Detection on page 1669](#).

You can enable or disable automatic re-keying between IKE peers through the `phase1-rekey` attribute of the `config system global` CLI command. For more information, see the “System” chapter of the [FortiGate CLI Reference](#).



When in FIPS-CC mode, the FortiGate unit requires DH key exchange to use values at least 3072 bits long. However most browsers need the key size set to 1024. You can set the minimum size of the DH keys in the CLI.

```
config system global
    set dh-params 3072
end
```

When you use a preshared key (shared secret) to set up two-party authentication, the remote VPN peer or client and the FortiGate unit must both be configured with the same preshared key. Each party uses a session key derived from the Diffie-Hellman exchange to create an authentication key, which is used to sign a known combination of inputs using an authentication algorithm (such as HMAC-MD5, HMAC-SHA-1, or HMAC-SHA-256). Hash-based Message Authentication Code (HMAC) is a method for calculating an authentication code

using a hash function plus a secret key, and is defined in [RFC 2104](#). Each party signs a different combination of inputs and the other party verifies that the same result can be computed.



For information regarding NP accelerated offloading of IPsec VPN authentication algorithms, please refer to the [Hardware Acceleration](#) handbook chapter.

When you use preshared keys to authenticate VPN peers or clients, you must distribute matching information to all VPN peers and/or clients whenever the preshared key changes.

As an alternative, the remote peer or dialup client and FortiGate unit can exchange digital signatures to validate each other's identity with respect to their public keys. In this case, the required digital certificates must be installed on the remote peer and on the FortiGate unit. By exchanging certificate DNs, the signed server certificate on one peer is validated by the presence of the root certificate installed on the other peer.

The following procedure assumes that you already have a Phase 1 definition that describes how remote VPN peers and clients will be authenticated when they attempt to connect to a local FortiGate unit. For information about the Local ID and XAuth options, see [Defining IKE negotiation parameters on page 1666](#) and [Defining IKE negotiation parameters on page 1666](#). Follow this procedure to add IKE negotiation parameters to the existing definition.

Defining IKE negotiation parameters

1. Go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
2. Edit the **Phase 1 Proposal** (if it is not available, you may need to click the **Convert to Custom Tunnel** button).
3. Select **Phase 1 Proposal** and include the appropriate entries as follows:

Phase 1 Proposal

Select the encryption and authentication algorithms that will be used to generate keys for protecting negotiations.

Add or delete encryption and authentication algorithms as required. Select a minimum of one and a maximum of three combinations. The remote peer must be configured to use at least one of the proposals that you define.

It is invalid to set both **Encryption** and **Authentication** to null.

Encryption

Select a symmetric-key algorithms:

NULL — Do not use an encryption algorithm.

DES — Digital Encryption Standard, a 64-bit block algorithm that uses a 56-bit key.

3DES — Triple-DES; plain text is encrypted three times by three keys.

AES128 — A 128-bit block algorithm that uses a 128-bit key.

AES192 — A 128-bit block algorithm that uses a 192-bit key.

AES256 — A 128-bit block algorithm that uses a 256-bit key.

ChaCha20/Poly1305 — A 128-bit block algorithm that uses a 128-bit key and a symmetric cipher. Only available for IKEv2.

Authentication	<p>You can select either of the following message digests to check the authenticity of messages during an encrypted session:</p> <p>NULL — Do not use a message digest. MD5 — Message Digest 5. SHA1 — Secure Hash Algorithm 1 - a 160-bit message digest.</p> <p>To specify one combination only, set the Encryption and Authentication options of the second combination to NULL. To specify a third combination, use the Add button beside the fields for the second combination.</p> <p>For information regarding NP accelerated offloading of IPsec VPN authentication algorithms, please refer to the Hardware Acceleration handbook chapter.</p>
Diffie-Hellman Group	<p>Select one or more Diffie-Hellman groups from DH groups 1, 2, 5, 14 through 21, and 27 through 30. When using aggressive mode, DH groups cannot be negotiated. By default, DH group 14 is selected, to provide sufficient protection for stronger cipher suites that include AES and SHA2. If you select multiple DH groups, the order they appear in the configuration is the order in which they are negotiates.</p> <p>If both VPN peers (or a VPN server and its client) have static IP addresses and use aggressive mode, select a single DH group. The setting on the FortiGate unit must be identical to the setting on the remote peer or dialup client.</p> <p>When the remote VPN peer or client has a dynamic IP address and uses aggressive mode, select up to three DH groups on the FortiGate unit and one DH group on the remote peer or dialup client. The setting on the remote peer or dialup client must be identical to one of the selections on the FortiGate unit.</p> <p>If the VPN peer or client employs main mode, you can select multiple DH groups. At least one of the settings on the remote peer or dialup client must be identical to the selections on the FortiGate unit.</p>
Keylife	<p>Type the amount of time (in seconds) that will be allowed to pass before the IKE encryption key expires. When the key expires, a new key is generated without interrupting service. The keylife can be from 120 to 172800 seconds.</p>
Nat-traversal	<p>Enable this option if a NAT device exists between the local FortiGate unit and the VPN peer or client. The local FortiGate unit and the VPN peer or client must have the same NAT traversal setting (both selected or both cleared). When in doubt, enable NAT-traversal. See NAT traversal on page 1669.</p>

Keepalive Frequency	If you enabled NAT traversal, enter a keepalive frequency setting. The value represents an interval from 0 to 900 seconds where the connection will be maintained with no activity. For additional security this value must be as low as possible. See NAT keepalive frequency on page 1669 .
Dead Peer Detection	Enable this option to reestablish VPN tunnels on idle connections and clean up dead IKE peers if required. This feature minimizes the traffic required to check if a VPN peer is available or unavailable (dead). See Dead Peer Detection on page 1669 .

NAT traversal

Network Address Translation (NAT) is a way to convert private IP addresses to publicly routable Internet addresses and vice versa. When an IP packet passes through a NAT device, the source or destination address in the IP header is modified. FortiGate units support NAT version 1 (encapsulate on port 500 with non-IKE marker), version 3 (encapsulate on port 4500 with non-ESP marker), and compatible versions.

NAT cannot be performed on IPsec packets in ESP tunnel mode because the packets do not contain a port number. As a result, the packets cannot be demultiplexed. To work around this, the FortiGate unit provides a way to protect IPsec packet headers from NAT modifications. When the Nat-traversal option is enabled, outbound encrypted packets are wrapped inside a UDP IP header that contains a port number. This extra encapsulation allows NAT devices to change the port number without modifying the IPsec packet directly.

To provide the extra layer of encapsulation on IPsec packets, the Nat-traversal option must be enabled whenever a NAT device exists between two FortiGate VPN peers or a FortiGate unit and a dialup client such as FortiClient. On the receiving end, the FortiGate unit or FortiClient removes the extra layer of encapsulation before decrypting the packet.

Additionally, you can force IPsec to use NAT traversal. If NAT is set to **Forced**, the FortiGate will use a port value of zero when constructing the NAT discovery hash for the peer. This causes the peer to think it is behind a NAT device, and it will use UDP encapsulation for IPsec, even if no NAT is present. This approach maintains interoperability with any IPsec implementation that supports the NAT-T RFC.

NAT keepalive frequency

When a NAT device performs network address translation on a flow of packets, the NAT device determines how long the new address will remain valid if the flow of traffic stops (for example, the connected VPN peer may be idle). The device may reclaim and reuse a NAT address when a connection remains idle for too long.

To work around this, when you enable NAT traversal specify how often the FortiGate unit sends periodic keepalive packets through the NAT device in order to ensure that the NAT address mapping does not change during the lifetime of a session. To be effective, the keepalive interval must be smaller than the session lifetime value used by the NAT device.

The keepalive packet is a 138-byte ISAKMP exchange.

Dead Peer Detection

Sometimes, due to routing issues or other difficulties, the communication link between a FortiGate unit and a VPN peer or client may go down. Packets could be lost if the connection is left to time out on its own. The FortiGate unit provides a mechanism called Dead Peer Detection (DPD), sometimes referred to as gateway detection or ping server, to prevent this situation and reestablish IKE negotiations automatically before a

connection times out: the active Phase 1 security associations are caught and renegotiated (rekeyed) before the Phase 1 encryption key expires.

By default, Dead Peer Detection sends probe messages every five seconds by default (see `dpd-retryinterval` in the [FortiGate CLI Reference](#)). If you are experiencing high network traffic, you can experiment with increasing the ping interval. However longer intervals will require more traffic to detect dead peers which will result in more traffic.

In the web-based manager, the Dead Peer Detection option can be enabled when you define advanced Phase 1 options. The `config vpn ipsec phase1` CLI command supports additional options for specifying a retry count and a retry interval.

For more information about these commands and the related `config router gwdetect` CLI command, see the [FortiGate CLI Reference](#).

For example, enter the following CLI commands to configure dead peer detection on the existing IPsec Phase 1 configuration called `test` to use 15 second intervals and to wait for 3 missed attempts before declaring the peer dead and taking action.

```
config vpn ipsec phase1-interface
  edit <value>
    set dpd [disable | on-idle | on-demand]
    set dpd-retryinterval 15
    set dpd-retrycount 3
  next
end
```



The default for `vpn ipsec phase1 dpd` is `on-idle` when the type is `dynamic` to encourage dialup server configurations to more pro-actively delete tunnels if the peer goes away.

DPD scalability

On a dial-up server, if a multitude of VPN connections are idle, the increased DPD exchange could negatively impact the performance/load of the daemon. For this reason, an option is available in the CLI to send DPD passively in a mode called "on-demand".



- When there is no traffic and the last DPD-ACK had been received, IKE will not send DPDs periodically.
- IKE will only send out DPDs if there are outgoing packets to send but no inbound packets had since been received.

Syntax

Set DPD to `on-demand` to trigger DPD when IPsec traffic is sent but no reply is received from the peer.

```
config vpn ipsec phase1-interface
  edit <value>
    set dpd [disable | on-idle | on-demand]
  next
end
```

Certificate key size control

Proxy will choose the same SSL key size as the HTTPS server. If the key size from the server is 512, the proxy will choose 1024. If the key size is bigger than 1024, the proxy will choose 2048.

As a result, the `firewall ssl-ssh-profile` commands `certname-rsa`, `certname-dsa`, and `certname-ecdsa` have been replaced with more specific key size control commands under `vpn certificate` setting.

CLI syntax

```
config vpn certificate setting
  set certname-rsa1024 <name>
  set certname-rsa2048 <name>
  set certname-dsa1024 <name>
  set certname-dsa2048 <name>
  set certname-ecdsa256 <name>
  set certname-ecdsa384 <name>
end
```

Quantum resistant IKEv2 SA negotiation

An IKEv2 extension is available that changes the key generation mechanism to include a Post-quantum Pre-shared Key.

The addition of PPK in the calculation means that even if a quantum computer can break the Diffie-Hellman calculation to derive the DH-generated secret key, the inclusion of the PPK in the key generation algorithm means that the attacker is still unable to derive the keys used to authenticate the IKE SA negotiation (and so cannot impersonate either party in the negotiation) nor the keys used in negotiating an IPsec SA (or IKE SA).

Syntax

```
config vpn ipsec phase1-interface
  edit <name>
    set ike-version 2
    set type dynamic
    set ppk {disable | allow | require}
    set ppk-secret <ASCII string or hexadecimal encoded with a leading 0x>
    set ppk-identity <string>
  next
end
config user local
  edit <name>
    set type password
    set ppk-secret <ASCII string or hexadecimal encoded with a leading 0x>
  next
end
```

For troubleshooting, `diagnose vpn ike ga list` can indicate whether PPK was negotiated.

- The 'PPK' at the gateway level indicates whether PPK was negotiated during the initial IKE SA negotiation.
- The 'PPK' at the IKE SA level indicates whether PPK was negotiated on this IKE SA.
- The 'child' at the IKE SA level indicates whether the IKE SA is an initial IKE SA or whether it is a child IKE SA. The above has 'child: no' and so it is initial IKE SA.

Using XAuth authentication

Extended authentication (XAuth) increases security by requiring the remote dialup client user to authenticate in a separate exchange at the end of Phase 1. XAuth draws on existing FortiGate user group definitions and uses established authentication mechanisms such as PAP, CHAP, RADIUS, and LDAP to authenticate dialup clients. You can configure a FortiGate unit to function either as an XAuth server or an XAuth client. If the server or client is attempting a connection using XAuth and the other end is not using XAuth, the failed connection attempts that are logged will not specify XAuth as the reason.

Using the FortiGate unit as an XAuth server

A FortiGate unit can act as an XAuth server for dialup clients. When the Phase 1 negotiation completes, the FortiGate unit challenges the user for a user name and password. It then forwards the user's credentials to an external RADIUS or LDAP server for verification.

If the user records on the RADIUS server have suitably configured Framed-IP-Address fields, you can assign client virtual IP addresses by XAuth instead of from a DHCP address range. See [Assigning VIPs by RADIUS user group on page 1](#).

The authentication protocol to use for XAuth depends on the capabilities of the authentication server and the XAuth client:

- Select **PAP Server** whenever possible.
- You must select **PAP Server** for all implementations of LDAP and some implementations of Microsoft RADIUS.
- Select **Auto Server** when the authentication server supports **CHAP Server** but the XAuth client does not. The FortiGate unit will use PAP to communicate with the XAuth client and CHAP to communicate with the authentication server. You can also use **Auto Server** to allow multiple source interfaces to be defined in an IPsec/IKE policy

Before you begin, create user accounts and user groups to identify the dialup clients that need to access the network behind the FortiGate dialup server. If password protection will be provided through an external RADIUS or LDAP server, you must configure the FortiGate dialup server to forward authentication requests to the authentication server. For information about these topics, see the [FortiGate User Authentication Guide](#).

Authenticating a dialup user group using XAuth settings

1. At the FortiGate dialup server, go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
2. Select **Convert To Custom Tunnel**.
3. Edit **XAUTH**, select the **Type** setting, which determines the type of encryption method to use between the XAuth client, the FortiGate unit and the authentication server. Select one of the following options:
 - **Disabled** — Disables XAuth settings.
 - **PAP Server** — Password Authentication Protocol.
 - **CHAP Server** — Challenge-Handshake Authentication Protocol.
 - **Auto Server** — Use PAP between the XAuth client and the FortiGate unit, and CHAP between the FortiGate unit and the authentication server.
4. From the **User Group** list, select the user group that needs to access the private network behind the FortiGate unit. The group must be added to the FortiGate configuration before it can be selected here. For multiple user groups to be defined in the IPsec/IKE policy, select **Inherit from policy**.
4. Select **OK**.

5. Create as many policies as needed, specifying **Source User(s)** and **Destination Address**.
For example, one policy could have **user1** have access to **test_local_subnet_1**, while **user2** has access to **test_local_subnet_2**.



As of FortiOS 5.4.1, when XAuth settings are enabled, **Inherit from policy** is only available under **PAP Server** and **CHAP Server**, not **Auto Server**. Because of this, only one user group may be defined for **Auto Server**.

Using the FortiGate unit as an XAuth client

If the FortiGate unit acts as a dialup client, the remote peer, acting as an XAuth server, might require a username and password. You can configure the FortiGate unit as an XAuth client, with its own username and password, which it provides when challenged.

Configuring the FortiGate dialup client as an XAuth client

1. At the FortiGate dialup client, go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
2. Edit the **Phase 1 Proposal** (if it is not available, you may need to click the **Convert to Custom Tunnel** button).
3. Under **XAuth**, select **Enable as Client**.
4. In the **Username** field, type the FortiGate PAP, CHAP, RADIUS, or LDAP user name that the FortiGate XAuth server will compare to its records when the FortiGate XAuth client attempts to connect.
5. In the **Password** field, type the password to associate with the user name.
6. Select **OK**.

Dynamic IPsec route control

You can add a route to a peer destination selector by using the `add-route` option, which is available for all dynamic IPsec Phases 1 and 2, for both policy-based and route-based IPsec VPNs. This option was previously only available when `mode-cfg` was enabled in Phase 1.

The `add-route` option adds a route to the FortiGate unit's routing information base when the dynamic tunnel is negotiated. You can use the `distance` and `priority` options to set the distance and priority of this route. If this results in a route with the lowest distance, it is added to the FortiGate unit's forwarding information base.

You can also enable `add-route` in any policy-based or route-based Phase 2 configuration that is associated with a dynamic (dialup) Phase 1. In Phase 2, `add-route` can be enabled, disabled, or set to use the same route as Phase 1.

The `add-route` feature is enabled by default and is configured in the CLI.

Syntax

Phase 1

```
config vpn ipsec
  edit <name>
    set type dynamic
    set add-route {enable | disable}
  end
end
```


Phase 2

```
config vpn ipsec {phase2 | phase2-interface}
  edit <name>
    set add-route {phase1 | enable | disable}
  end
end
```

Blocking IPsec SA negotiation

For interface-based IPsec, IPsec SA negotiation blocking can only be removed if the peer offers a wildcard selector. If a wildcard selector is offered then the wildcard route will be added to the routing table with the distance/priority value configured in Phase 1 and, if that is the route with the lowest distance, it is installed into the forwarding information base.

In cases where this occurs, it is important to ensure that the distance value configured on Phase 1 is set appropriately.

Phase 2 parameters

This section describes the Phase 2 parameters that are required to establish communication through a VPN.

The following topics are included in this section:

Phase 2 settings

After IPsec VPN Phase 1 negotiations complete successfully, Phase 2 negotiation begins. Phase 2 parameters define the algorithms that the FortiGate unit can use to encrypt and transfer data for the remainder of the session. The basic Phase 2 settings associate IPsec Phase 2 parameters with a Phase 1 configuration.

When defining Phase 2 parameters, you can choose any set of Phase 1 parameters to set up a secure connection and authenticate the remote peer.

For more information on Phase 2 settings in the web-based manager, see [IPsec VPN in the web-based manager on page 1642](#).

The information and procedures in this section do not apply to VPN peers that perform negotiations using manual keys.

Phase 2 proposals

In Phase 2, the VPN peer or client and the FortiGate unit exchange keys again to establish a secure communication channel. The Phase 2 Proposal parameters select the encryption and authentication algorithms needed to generate keys for protecting the implementation details of Security Associations (SAs). The keys are generated automatically using a Diffie-Hellman algorithm.

Replay detection

IPsec tunnels can be vulnerable to replay attacks. Replay Detection enables the FortiGate unit to check all IPsec packets to see if they have been received before. If any encrypted packets arrive out of order, the FortiGate unit discards them.

IKE/IPsec Extended Sequence Number (ESN) support

64-bit Extended Sequence numbers (as described in RFC 4303, RFC 4304 as an addition to IKEv1, and RFC 5996 for IKEv2.) are supported for IPsec when Replay Detection is enabled.

Perfect Forward Secrecy (PFS)

By default, Phase 2 keys are derived from the session key created in Phase 1. Perfect Forward Secrecy (PFS) forces a new Diffie-Hellman exchange when the tunnel starts and whenever the Phase 2 keylife expires, causing a new key to be generated each time. This exchange ensures that the keys created in Phase 2 are unrelated to the Phase 1 keys or any other keys generated automatically in Phase 2.

Keylife

The Keylife setting sets a limit on the length of time that a Phase 2 key can be used. The default units are seconds. Alternatively, you can set a limit on the number of kilobytes (KB) of processed data, or both. If you select

both, the key expires when either the time has passed or the number of KB have been processed. When the Phase 2 key expires, a new key is generated without interrupting service.

Quick mode selectors

Quick mode selectors determine which IP addresses can perform IKE negotiations to establish a tunnel. By only allowing authorized IP addresses access to the VPN tunnel, the network is more secure.

The default settings are as broad as possible: any IP address or configured address object, using any protocol, on any port.



While the drop down menus for specifying an address also show address groups, the use of address groups may not be supported on a remote endpoint device that is not a FortiGate.

The address groups are at the bottom of the list to make it easy to distinguish between addresses and address groups.

When configuring Quick Mode selector **Source address** and **Destination address**, valid options include IPv4 and IPv6 single addresses, IPv4 subnet, or IPv6 subnet. For more information on IPv6 IPsec VPN, see [Overview of IPv6 IPsec support on page 1](#).

There are some configurations that require specific selectors:

- The VPN peer is a third-party device that uses specific phase2 selectors.
- The FortiGate unit connects as a dialup client to another FortiGate unit, in which case (usually) you must specify a source IP address, IP address range, or subnet. However, this is not required if you are using dynamic routing and mode-cfg.

With FortiOS VPNs, your network has multiple layers of security, with quick mode selectors being an important line of defence.

- Routes guide traffic from one IP address to another.
- Phase 1 and Phase 2 connection settings ensure there is a valid remote end point for the VPN tunnel that agrees on the encryption and parameters.
- Quick mode selectors allow IKE negotiations only for allowed peers.
- Security policies control which IP addresses can connect to the VPN.
- Security policies also control what protocols are allowed over the VPN along with any bandwidth limiting.



FortiOS is limited with IKEv2 selector matching. When using IKEv2 with a named traffic selector, no more than 32 subnets per traffic selector are added, since FortiOS doesn't fully implement the IKEv2 selector matching rules.

The workaround is to use multiple Phase 2s. If the configuration is FGT <-> FGT, then the better alternative is to just use 0.0.0.0 <-> 0.0.0.0 and use the firewall policy for enforcement.

Using the add-route option

Consider using the `add-route` option to add a route to a peer destination selector. Phase 2 includes the option of allowing the `add-route` to automatically match the settings in Phase 1. For more information, refer to [Phase 1 parameters on page 1655](#).

Syntax

Phase 2

```
config vpn ipsec {phase2 | phase2-interface}
    edit <name>
        set add-route {phase1 | enable | disable}
    end
end
```

Configuring the Phase 2 parameters

If you are creating a hub-and-spoke configuration or an Internet-browsing configuration, you may have already started defining some of the required Phase 2 parameters. If so, edit the existing definition to complete the configuration.

Specifying the Phase 2 parameters

1. Go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
2. Open the **Phase 2 Selectors** panel (if it is not available, you may need to click the **Convert to Custom Tunnel** button).
3. Enter a **Name** for the Phase 2 configuration, and select a **Phase 1** configuration from the drop-down list.
4. Select **Advanced**.
5. Include the appropriate entries as follows:

Phase 2 Proposal

Select the encryption and authentication algorithms that will be used to change data into encrypted code.

Add or delete encryption and authentication algorithms as required. Select a minimum of one and a maximum of three combinations. The remote peer must be configured to use at least one of the proposals that you define.

It is invalid to set both **Encryption** and **Authentication** to null.

Encryption

Select a symmetric-key algorithms:

NULL — Do not use an encryption algorithm.

DES — Digital Encryption Standard, a 64-bit block algorithm that uses a 56-bit key.

3DES — Triple-DES; plain text is encrypted three times by three keys.

AES128 — A 128-bit block algorithm that uses a 128-bit key.

AES192 — A 128-bit block algorithm that uses a 192-bit key.

AES256 — A 128-bit block algorithm that uses a 256-bit key.

ChaCha20/Poly1305 — A 128-bit block algorithm that uses a 128-bit key and a symmetric cipher. Only available for IKEv2.

Authentication	<p>You can select either of the following message digests to check the authenticity of messages during an encrypted session:</p> <p>NULL — Do not use a message digest. MD5 — Message Digest 5. SHA1 — Secure Hash Algorithm 1 - a 160-bit message digest.</p> <p>To specify one combination only, set the Encryption and Authentication options of the second combination to NULL. To specify a third combination, use the Add button beside the fields for the second combination.</p> <p>For information regarding NP accelerated offloading of IPsec VPN authentication algorithms, please refer to the Hardware Acceleration handbook chapter.</p>
Enable replay detection	Optionally enable or disable replay detection. Replay attacks occur when an unauthorized party intercepts a series of IPsec packets and replays them back into the tunnel.
Enable perfect forward secrecy (PFS)	Enable or disable PFS. Perfect forward secrecy (PFS) improves security by forcing a new Diffie-Hellman exchange whenever keylife expires.
Diffie-Hellman Group	Select one Diffie-Hellman group (1, 2, 5, 14 through 21, or 27 through 30). The remote peer or dialup client must be configured to use the same group.
Keylife	Select the method for determining when the Phase 2 key expires: Seconds , KBytes , or Both . If you select Both , the key expires when either the time has passed or the number of KB have been processed. The range is from 120 to 172800 seconds, or from 5120 to 2147483648 KB.
Autokey Keep Alive	Enable the option if you want the tunnel to remain active when no data is being processed.
Auto-negotiate	Enable the option if you want the tunnel to be automatically renegotiated when the tunnel expires.
DHCP-IPsec	<p>Select Enable if the FortiGate unit acts as a dialup server and FortiGate DHCP server or relay will be used to assign VIP addresses to FortiClient dialup clients. The DHCP server or relay parameters must be configured separately.</p> <p>If the FortiGate unit acts as a dialup server and the FortiClient dialup client VIP addresses match the network behind the dialup server, select Enable to cause the FortiGate unit to act as a proxy for the dialup clients.</p> <p>This is available only for Phase 2 configurations associated with a dialup Phase 1 configuration. It works only on policy-based VPNs.</p>

Autokey Keep Alive

The Phase 2 SA has a fixed duration. If there is traffic on the VPN as the SA nears expiry, a new SA is negotiated and the VPN switches to the new SA without interruption. If there is no traffic, however, the SA expires (by default) and the VPN tunnel goes down. A new SA will not be generated until there is traffic.

The Autokey Keep Alive option ensures that a new Phase 2 SA is negotiated, even if there is no traffic, so that the VPN tunnel stays up.

Auto-negotiate

By default, the Phase 2 security association (SA) is not negotiated until a peer attempts to send data. The triggering packet and some subsequent packets are dropped until the SA is established. Applications normally resend this data, so there is no loss, but there might be a noticeable delay in response to the user.

If the tunnel goes down, the auto-negotiate feature (when enabled) attempts to re-establish the tunnel. Auto-negotiate initiates the Phase 2 SA negotiation automatically, repeating every five seconds until the SA is established.

Automatically establishing the SA can be important for a dialup peer. It ensures that the VPN tunnel is available for peers at the server end to initiate traffic to the dialup peer. Otherwise, the VPN tunnel does not exist until the dialup peer initiates traffic.

The auto-negotiate feature is available through the Command Line Interface (CLI) via the following commands:

```
config vpn ipsec phase2
  edit <phase2_name>
    set auto-negotiate enable
  end
```

Installing dynamic selectors via auto-negotiate

The IPsec SA connect message generated is used to install dynamic selectors. These selectors can now be installed via the auto-negotiate mechanism. When phase 2 has auto-negotiate enabled, and phase 1 has mesh-selector-type set to **subnet**, a new dynamic selector will be installed for each combination of source and destination subnets. Each dynamic selector will inherit the auto-negotiate option from the template selector and begin SA negotiation. Phase 2 selector sources from dial-up clients will all establish SAs without traffic being initiated from the client subnets to the hub.

DHCP-IPsec

Select this option if the FortiGate unit assigns VIP addresses to FortiClient dialup clients through a DHCP server or relay. This option is available only if the Remote Gateway in the Phase 1 configuration is set to Dialup User and it works only on policy-based VPNs.

With the DHCP-IPsec option, the FortiGate dialup server acts as a proxy for FortiClient dialup clients that have VIP addresses on the subnet of the private network behind the FortiGate unit. In this case, the FortiGate dialup server acts as a proxy on the local private network for the FortiClient dialup client. A host on the network behind the dialup server issues an ARP request, corresponding to the device MAC address of the FortiClient host (when a remote server sends an ARP to the local FortiClient dialup client). The FortiGate unit then answers the ARP request on behalf of the FortiClient host, and forwards the associated traffic to the FortiClient host through the tunnel.

Acting as a proxy prevents the VIP address assigned to the FortiClient dialup client from causing possible ARP broadcast problems — the normal and VIP addresses can confuse some network switches by two addresses having the same MAC address.

IPsec support for ChaCha20/Poly1305 AEAD cipher

In IKEv2, to support [RFC 7634](#), crypto algorithms ChaCha20 and Poly1305 can be used together as a combined mode AEAD cipher (like aes-gcm) in the new `crypto_ftnt cipher` in `cipher_chacha20poly1305.c`.

Syntax

```
config vpn ipsec phase2-interface
  edit <name>
    set phase1name <name>
    set proposal chacha20poly1305
  next
end
```

IPsec support for AES-GCM for IKEv2 Phase 1

In IKEv2, to support [RFC 5282](#), AEAD algorithm AES-GCM is now supported, both 128 and 256-bit variants.

Syntax

```
config vpn ipsec phase2-interface
  edit <name>
    set phase1name <name>
    set proposal [aes128gcm | aes256gcm]
  next
end
```

Defining VPN security policies

This section explains how to specify the source and destination IP addresses of traffic transmitted through an IPsec VPN, and how to define appropriate security policies.

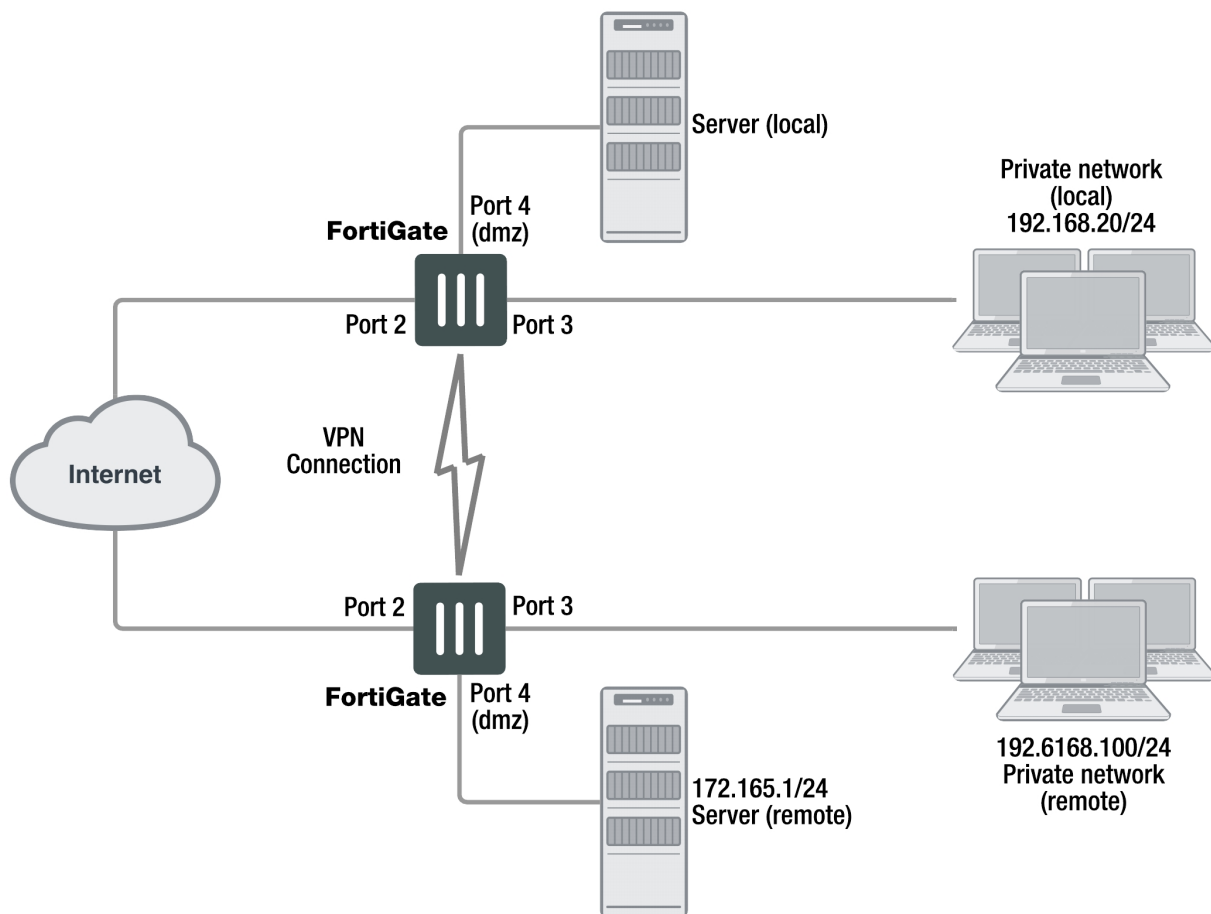
The following topics are included in this section:

Defining policy addresses

A VPN tunnel has two end points. These end points may be VPN peers such as two FortiGate gateways. Encrypted packets are transmitted between the end points. At each end of the VPN tunnel, a VPN peer intercepts encrypted packets, decrypts the packets, and forwards the decrypted IP packets to the intended destination.

You need to define firewall addresses for the private networks behind each peer. You will use these addresses as the source or destination address depending on the security policy.

Example topology for the following policies



In general:

- In a gateway-to-gateway, hub-and-spoke, dynamic DNS, redundant-tunnel, or transparent configuration, you need to define a policy address for the private IP address of the network behind the remote VPN peer (for example, 192.168.10.0/255.255.255.0 or 192.168.10.0/24).
- In a peer-to-peer configuration, you need to define a policy address for the private IP address of a server or host behind the remote VPN peer (for example, 172.16.5.1/255.255.255.255 or 172.16.5.1/32 or 172.16.5.1).

For a FortiGate dialup server in a dialup-client or Internet-browsing configuration:

- If you are not using VIP addresses, or if the FortiGate dialup server assigns VIP addresses to FortiClient dialup clients through FortiGate DHCP relay, select the predefined destination address “all” in the security policy to refer to the dialup clients.
- If you assign VIP addresses to FortiClient dialup clients manually, you need to define a policy address for the VIP address assigned to the dialup client (for example, 10.254.254.1/32), or a subnet address from which the VIP addresses are assigned (for example, 10.254.254.0/24 or 10.254.254.0/255.255.255.0).
- For a FortiGate dialup client in a dialup-client or Internet-browsing configuration, you need to define a policy address for the private IP address of a host, server, or network behind the FortiGate dialup server.

Defining a security IP address

1. Go to **Policy & Objects > Addresses** and select **Create New**.
2. In the **Name** field, type a descriptive name that represents the network, server(s), or host(s).
3. In **Type**, select **Subnet**.
4. In the **Subnet/IP Range** field, type the corresponding IP address and subnet mask.
For a subnet you could use the format 172.16.5.0/24 or its equivalent 172.16.5.0/255.255.255.0. For a server or host it would likely be 172.16.5.1/32. Alternately you can use an IP address range such as 192.168.10.[80-100] or 192.168.10.80-192.168.10.100.
5. Select **OK**.

Defining security policies for policy-based and route-based VPNs

Security policies allow IP traffic to pass between interfaces on a FortiGate unit. You can limit communication to particular traffic by specifying source address and destination addresses. Then only traffic from those addresses will be allowed.

Policy-based and route-based VPNs require different security policies.

- A policy-based VPN requires an IPsec security policy. You specify the interface to the private network, the interface to the remote peer and the VPN tunnel. A single policy can enable traffic inbound, outbound, or in both directions.
- A route-based VPN requires an Accept security policy for each direction. As source and destination interfaces, you specify the interface to the private network and the virtual IPsec interface (Phase 1 configuration) of the VPN. The IPsec interface is the destination interface for the outbound policy and the source interface for the inbound policy. One security policy must be configured for each direction of each VPN interface.

There are examples of security policies for both policy-based and route-based VPNs throughout this guide. See [Route-based or policy-based VPN on page 1725](#).



If the security policy, which grants the VPN Connection is limited to certain services, DHCP must be included, otherwise the client won't be able to retrieve a lease from the FortiGate's (IPsec) DHCP server, because the DHCP Request (coming out of the tunnel) will be blocked.

Policy-based VPN

An IPsec security policy enables the transmission and reception of encrypted packets, specifies the permitted direction of VPN traffic, and selects the VPN tunnel. In most cases, a single policy is needed to control both inbound and outbound IP traffic through a VPN tunnel. Be aware of the following considerations below before creating an IPsec security policy.

Allow traffic to be initiated from the remote site

Security policies specify which IP addresses can initiate a tunnel. By default, traffic from the local private network initiates the tunnel. When the **Allow traffic to be initiated from the remote site** option is selected, traffic from a dialup client, or a computer on a remote network, initiates the tunnel. Both can be enabled at the same time for bi-directional initiation of the tunnel.

Outbound and inbound NAT

When a FortiGate unit operates in NAT mode, you can also enable inbound or outbound NAT. Outbound NAT may be performed on outbound encrypted packets or IP packets in order to change their source address before they are sent through the tunnel. Inbound NAT is performed to intercept and decrypt emerging IP packets from the tunnel.

By default, these options are not selected in security policies and can only be set through the CLI. For more information on this, see the “config firewall” chapter of the [FortiGate CLI Reference](#).

Source and destination addresses

Most security policies control outbound IP traffic. A VPN outbound policy usually has a source address originating on the private network behind the local FortiGate unit, and a destination address belonging to a dialup VPN client or a network behind the remote VPN peer. The source address that you choose for the security policy identifies from where outbound cleartext IP packets may originate, and also defines the local IP address or addresses that a remote server or client will be allowed to access through the VPN tunnel. The destination address that you choose identifies where IP packets must be forwarded after they are decrypted at the far end of the tunnel, and determines the IP address or addresses that the local network will be able to access at the far end of the tunnel.

Enabling other policy features

You can fine-tune a policy for services such as HTTP, FTP, and POP3, enable logging, traffic shaping, antivirus protection, web filtering, email filtering, file transfer, email services, and optionally allow connections according to a predefined schedule.

As an option, differentiated services (diffserv or DSCP) for the security policy can be enabled through the CLI. For more information on this feature, see the [Traffic Shaping](#) handbook chapter, or the “firewall” chapter of the [FortiGate CLI Reference](#).

Before you begin

Before you define the IPsec policy, you must:

- Define the IP source and destination addresses. See [Defining policy addresses on page 1681](#).
- Specify the Phase 1 authentication parameters. See [Phase 1 parameters on page 1655](#).
- Specify the Phase 2 parameters. See [Phase 2 parameters on page 1675](#).

Defining an IPsec security policy

1. Go to **Policy & Objects > IPv4 Policy**.
2. Select **Create New** and set the following options:

Name	Enter a name for the security policy.
Incoming Interface	Select the local interface to the internal (private) network.
Outgoing Interface	Select the local interface to the external (public) network.
Source	Select the name that corresponds to the local network, server(s), or host(s) from which IP packets may originate.
Destination Address	Select the name that corresponds to the remote network, server(s), or host(s) to which IP packets may be delivered.
Schedule	Keep the default setting (always) unless changes are needed to meet specific requirements.
Service	Keep the default setting (ANY) unless changes are needed to meet your specific requirements.
Action	For the purpose of this configuration, set Action to IPsec . Doing this will close Firewall / Network Options and open VPN Tunnel options. Select the VPN tunnel of your choice, and select Allow traffic to be initiated from the remote site , which will allow traffic from the remote network to initiate the tunnel.

3. You may enable UTM features, and/or event logging, or select advanced settings to authenticate a user group, or shape traffic. For more information, see the [Firewall](#) handbook chapter.
4. Select **OK**.
5. Place the policy in the policy list above any other policies having similar source and destination addresses.

Defining multiple IPsec policies for the same tunnel

You must define at least one IPsec policy for each VPN tunnel. If the same remote server or client requires access to more than one network behind a local FortiGate unit, the FortiGate unit must be configured with an IPsec policy for each network. Multiple policies may be required to configure redundant connections to a remote destination or control access to different services at different times.

To ensure a secure connection, the FortiGate unit must evaluate policies with **Action** set to **IPsec** before **ACCEPT** and **DENY**. Because the FortiGate unit reads policies starting at the top of the list, you must move all IPsec policies to the top of the list, and be sure to reorder your multiple IPsec policies that apply to the tunnel so that specific constraints can be evaluated before general constraints.



Adding multiple IPsec policies for the same VPN tunnel can cause conflicts if the policies specify similar source and destination addresses, but have different settings for the same service. When policies overlap in this manner, the system may apply the wrong IPsec policy or the tunnel may fail.

For example, if you create two equivalent IPsec policies for two different tunnels, it does not matter which one comes first in the list of IPsec policies — the system will select the correct policy based on the specified source and destination addresses. If you create two different IPsec policies for the same tunnel (that is, the two policies treat traffic differently depending on the nature of the connection request), you might have to reorder the IPsec policies to ensure that the system selects the correct IPsec policy.

Route-based VPN

When you define a route-based VPN, you create a virtual IPsec interface on the physical interface that connects to the remote peer. You create ordinary Accept security policies to enable traffic between the IPsec interface and the interface that connects to the private network. This makes configuration simpler than for policy-based VPNs, which require IPsec security policies.

Defining security policies for a route-based VPN

1. Go to **Policy & Objects > IPv4 Policy**.
2. Select **Create New** and define an **ACCEPT** security policy to permit communication between the local private network and the private network behind the remote peer. Enter these settings in particular:

Name	Enter a name for the security policy.
Incoming Interface	Select the interface that connects to the private network behind this FortiGate unit.
Outgoing Interface	Select the IPsec Interface you configured.
Source	Select the address name that you defined for the private network behind this FortiGate unit.
Destination Address	Select the address name that you defined for the private network behind the remote peer.
Action	Select ACCEPT .
NAT	Disable NAT .

To permit the remote client to initiate communication, you need to define a security policy for communication in that direction.

3. Select **Create New** and enter these settings in particular:

Name	Enter a name for the security policy.
Incoming Interface	Select the IPsec Interface you configured.
Outgoing Interface	Select the interface that connects to the private network behind this FortiGate unit.
Source	Select the address name that you defined for the private network behind the remote peer.

Destination Address	Select the address name that you defined for the private network behind this FortiGate unit.
Action	Select ACCEPT .
NAT	Disable NAT .

Gateway-to-gateway

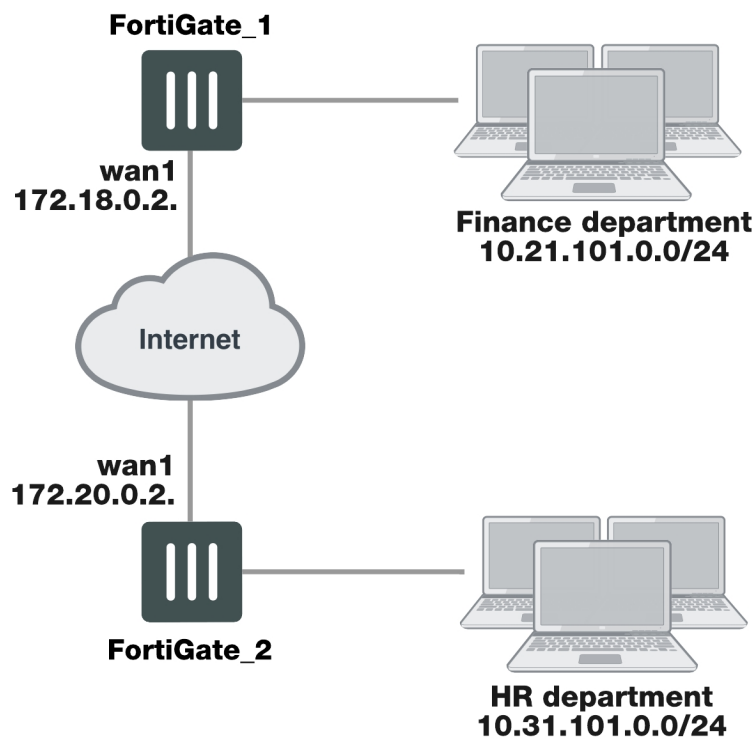
This section explains how to set up a basic gateway-to-gateway (site-to-site) IPsec VPN.

The following topics are included in this section:

Configuration overview

In a gateway-to-gateway configuration, two FortiGate units create a VPN tunnel between two separate private networks. All traffic between the two networks is encrypted and protected by FortiGate security policies.

Example gateway-to-gateway configuration




In some cases, computers on the private network behind one VPN peer may (by co-incidence) have IP addresses that are already used by computers on the network behind the other VPN peer. In this type of situation (ambiguous routing), conflicts may occur in one or both of the FortiGate routing tables and traffic destined for the remote network through the tunnel may not be sent. To resolve issues related to ambiguous routing, see [Configuration overview on page 1687](#).

In other cases, computers on the private network behind one VPN peer may obtain IP addresses from a local DHCP server. However, unless the local and remote networks use different private network address spaces, unintended ambiguous routing and/or IP-address overlap issues may arise. For a discussion of the related issues, see [FortiGate dialup-client configurations on page 1](#).

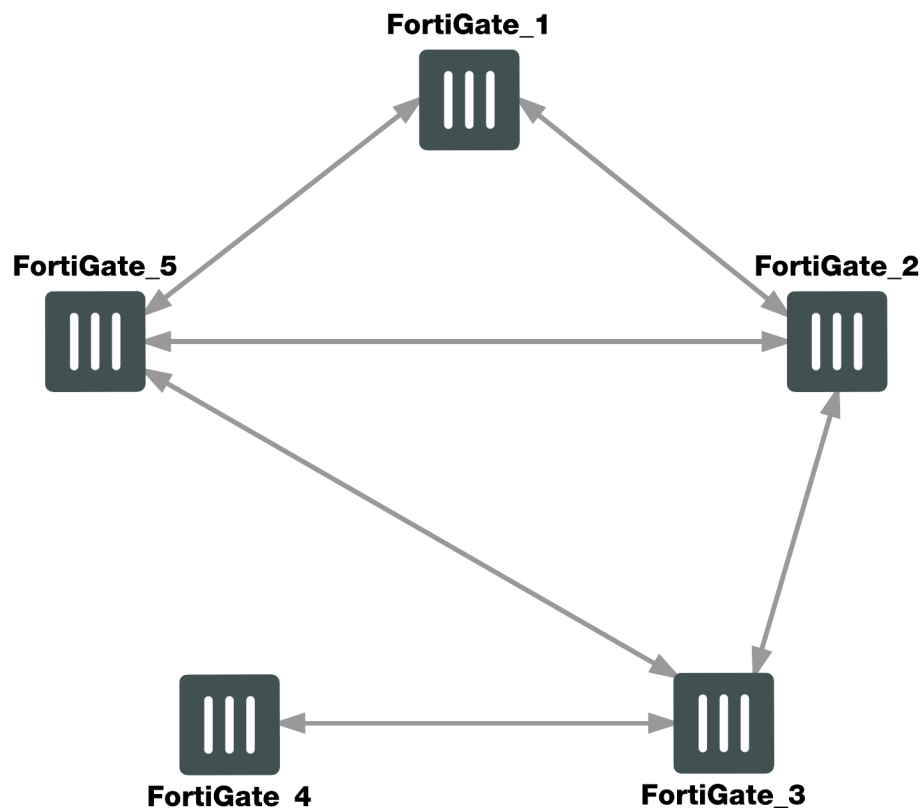
You can set up a fully meshed or partially meshed configuration (see below).

Fully meshed configuration



In a fully meshed network, all VPN peers are connected to each other, with one hop between peers. This topology is the most fault-tolerant: if one peer goes down, the rest of the network is not affected. This topology is difficult to scale because it requires connections between all peers. In addition, unnecessary communication can occur between peers. Best practices dictates a hub-and-spoke configuration instead (see [Hub-and-spoke configurations on page 1](#)).

Partially meshed configuration



A partially meshed network is similar to a fully meshed network, but instead of having tunnels between all peers, tunnels are only configured between peers that communicate with each other regularly.

Gateway-to-gateway configuration

The FortiGate units at both ends of the tunnel must be operating in NAT mode and have static public IP addresses.

When a FortiGate unit receives a connection request from a remote VPN peer, it uses IPsec Phase 1 parameters to establish a secure connection and authenticate that VPN peer. Then, if the security policy permits the connection, the FortiGate unit establishes the tunnel using IPsec Phase 2 parameters and applies the IPsec security policy. Key management, authentication, and security services are negotiated dynamically through the IKE protocol.

To support these functions, the following general configuration steps must be performed by both FortiGate units:

- Define the Phase 1 parameters that the FortiGate unit needs to authenticate the remote peer and establish a secure connection.
- Define the Phase 2 parameters that the FortiGate unit needs to create a VPN tunnel with the remote peer.
- Create security policies to control the permitted services and permitted direction of traffic between the IP source and destination addresses.

Configuring Phase 1 and Phase 2 for both peers

This procedure applies to both peers. Repeat the procedure on each FortiGate unit, using the correct IP address for each. You may wish to vary the Phase 1 names but this is optional. Otherwise all steps are the same for each peer.

The Phase 1 configuration defines the parameters that FortiGate_1 will use to authenticate FortiGate_2 and establish a secure connection. For the purposes of this example, a preshared key will be used to authenticate FortiGate_2. The same preshared key must be specified at both FortiGate units.

Before you define the Phase 1 parameters, you need to:

- Reserve a name for the remote gateway.
- Obtain the IP address of the public interface to the remote peer.
- Reserve a unique value for the preshared key.

The key must contain at least 6 printable characters and best practices dictate that it only be known by network administrators. For optimum protection against currently known attacks, the key must have a minimum of 16 randomly chosen alphanumeric characters.

At the local FortiGate unit, define the Phase 1 configuration needed to establish a secure connection with the remote peer. See [IPsec VPN in the web-based manager on page 1642](#).

1. Go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
2. Edit the **Phase 1 Proposal** (if it is not available, you may need to click the **Convert to Custom Tunnel** button).
3. Enter the following information, and select **OK**.

Name	Enter <code>peer_1</code> . A name to identify the VPN tunnel. This name appears in Phase 2 configurations, security policies and the VPN monitor.
Remote Gateway	Select Static IP Address .
IP Address	Enter <code>172.20.0.2</code> when configuring FortiGate_1. Enter <code>172.18.0.2</code> when configuring FortiGate_2. The IP address of the remote peer public interface.
Local Interface	Select wan1 .

The basic Phase 2 settings associate IPsec Phase 2 parameters with the Phase 1 configuration and specify the remote end point of the VPN tunnel. Before you define the Phase 2 parameters, you need to reserve a name for the tunnel. See [IPsec VPN in the web-based manager on page 1642](#).

1. Open the **Phase 2 Selectors** panel (if it is not available, you may need to click the **Convert to Custom Tunnel** button).
2. Enter a **Name** of `peer_1_p2`.
3. Select **peer_1** from the **Phase 1** drop-down menu.

Creating security policies

Security policies control all IP traffic passing between a source address and a destination address.

An IPsec security policy is needed to allow the transmission of encrypted packets, specify the permitted direction of VPN traffic, and select the VPN tunnel that will be subject to the policy. A single policy is needed to control both inbound and outbound IP traffic through a VPN tunnel.

Before you define security policies, you must first specify the IP source and destination addresses. In a gateway-to-gateway configuration:

- The IP source address corresponds to the private network behind the local FortiGate unit.
- The IP destination address refers to the private network behind the remote VPN peer.

When you are creating security policies, choose one of either route-based or policy-based methods and follow it for both VPN peers. DO NOT configure both route-based and policy-based policies on the same FortiGate unit for the same VPN tunnel.

The configuration of FortiGate_2 is similar to that of FortiGate_1. You must:

- Define the Phase 1 parameters that FortiGate_2 needs to authenticate FortiGate_1 and establish a secure connection.
- Define the Phase 2 parameters that FortiGate_2 needs to create a VPN tunnel with FortiGate_1.
- Create the security policy and define the scope of permitted services between the IP source and destination addresses.

When creating security policies it is good practice to include a comment describing what the policy does.

Creating firewall addresses

Define names for the addresses or address ranges of the private networks that the VPN links. These addresses are used in the security policies that permit communication between the networks.

To define the IP address of the network behind FortiGate_1

1. Go to **Policy & Objects > Addresses** and select **Create New**.
2. Enter the **Name** of `Finance_network`.
3. Select a **Type** of **Subnet**.
4. Enter the **Subnet** of `10.21.101.0/24`.
5. Select **OK**.

To specify the address of the network behind FortiGate_2

1. Go to **Policy & Objects > Addresses** and select **Create New**.
2. Enter the **Name** of `HR_network`.
3. Select a **Type** of **Subnet**.
4. Enter the **Subnet/IP Range** of `10.31.101.0/24`.
5. Select **OK**.

Creating route-based VPN security policies

Define an ACCEPT security policy to permit communications between the source and destination addresses.

To create route-based VPN security policies

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Leave the **Policy Type** as **Firewall** and leave the **Policy Subtype** as **Address**.

3. Enter the following, and select **OK**.

Incoming Interface	Select internal . The interface that connects to the private network behind this FortiGate unit.
Source Address	Select Finance_network when configuring FortiGate_1. Select HR_network when configuring FortiGate_2. The address name for the private network behind this FortiGate unit.
Outgoing Interface	Select peer_1 . The VPN Tunnel (IPsec Interface) you configured earlier.
Destination Address	Select HR_network when configuring FortiGate_1. Select Finance_network when configuring FortiGate_2. The address name that you defined for the private network behind the remote peer.
Action	Select ACCEPT .
Enable NAT	Disable.
Comments	Allow Internal to remote VPN network traffic.

4. Optionally, configure any additional features you may want, such as UTM or traffic shaping.
 5. Select **Create New** to create another policy for the other direction.
 6. Leave the **Policy Type** as **Firewall** and leave the **Policy Subtype** as **Address**.
 7. Enter the following information, and select **OK**.

Incoming Interface	Select peer_1 . The VPN Tunnel (IPsec Interface) you configured.
Source Address	Select HR_network when configuring FortiGate_1. Select Finance_Network when configuring FortiGate_2. The address name defined for the private network behind the remote peer.
Outgoing Interface	Select internal . The interface that connects to the private network behind this FortiGate unit.

Destination Address	Select Finance_Network when configuring FortiGate_1. Select HR_network when configuring FortiGate_2. The address name defined for the private network behind this FortiGate unit.
Action	Select ACCEPT .
Enable NAT	Disable.
Comments	Allow remote VPN network traffic to Internal.

8. Configure any additional features such as UTM or traffic shaping you may want. (optional).

All network traffic must have a static route to direct its traffic to the proper destination. Without a route, traffic will not flow even if the security policies are configured properly. You may need to create a static route entry for both directions of VPN traffic if your security policies allow bi-directional tunnel initiation.

To configure the route for a route-based VPN:

1. On FortiGate_2, go to **Network > Static Routes** and select **Create New**.
2. Enter the following information, and then select **OK**:

Destination IP / Mask	10.21.101.0/24
Device	FGT2_to_FGT1_Tunnel
Gateway	Leave as default: 0.0.0.0.
Distance (Advanced)	Leave this at its default. If there are other routes on this FortiGate unit, you may need to set the distance on this route so the VPN traffic will use it as the default route. However, this normally happens by default because this route is typically a better match than the generic default route.

Creating policy-based VPN security policy

Define an IPsec security policy to permit communications between the source and destination addresses.

1. Go to **Policy & Objects > IPv4 Policy**.
2. Complete the following:

Incoming Interface	Select internal . The interface that connects to the private network behind this FortiGate unit.
---------------------------	--

Source Address	Select Finance_network when configuring FortiGate_1. Select HR_network when configuring FortiGate_2. The address name defined for the private network behind this FortiGate unit.
Outgoing Interface	Select wan1 . The FortiGate unit's public interface.
Destination Address	Select HR_network when configuring FortiGate_1. Select Finance_network when configuring FortiGate_2.
VPN Tunnel	Select Use Existing and select peer_1 from the VPN Tunnel drop-down list. Select Allow traffic to be initiated from the remote site to enable traffic from the remote network to initiate the tunnel.
Comments	Bidirectional policy-based VPN policy.

Place VPN policies in the policy list above any other policies having similar source and destination addresses.

Remote Internet browsing for Site-to-Site VPN from the IPsec VPN Wizard

The IPsec VPN Wizard **Policy & Routing** section includes **Internet Access** options to support selecting **Share WAN** and **Force to use remote WAN**:

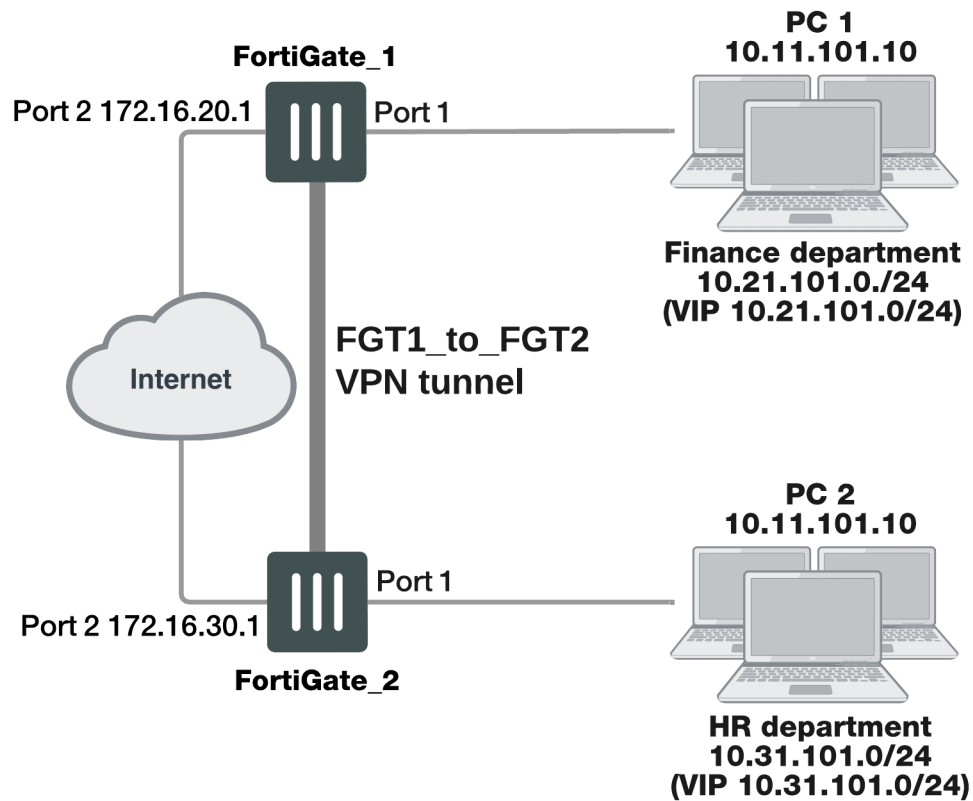
- The **Share WAN** option allows the remote subnet to browse the Internet via this FortiGate. When **Share WAN** is selected, a dropdown appears for the user to select the desired **Shared WAN**.
- The **Force to use remote WAN** option will send all Internet browsing traffic to the remote VPN gateway. The remote gateway must be configured with the **Share WAN** option enabled. When **Force to use remote WAN** is selected, a **Local Gateway** field appears (since a static route needs to be created to reach the remote gateway, because all other addresses will be routed via the VPN tunnel).

How to work with overlapping subnets

A site-to-site VPN configuration sometimes has the problem that the private subnet addresses at each end are the same. You can resolve this problem by remapping the private addresses using virtual IP addresses (VIP).

VIPs allow computers on those overlapping private subnets to each have another set of IP addresses that can be used without confusion. The FortiGate unit maps the VIP addresses to the original addresses. This means if PC1 starts a session with PC2 at 10.31.101.10, FortiGate_2 directs that session to 10.11.101.10 — the actual IP address of PC2. The figure below demonstrates this — Finance network VIP is 10.21.101.0/24 and the HR network is 10.31.101.0/24.

Overlapped subnets example



Solution for route-based VPN

You need to:

- Configure IPsec Phase 1 and Phase 2 as you usually would for a route-based VPN. In this example, the resulting IPsec interface is named `FGT1_to_FGT2`.
- Configure virtual IP (VIP) mapping:
 - the 10.21.101.0/24 network mapped to the 10.11.101.0/24 network on FortiGate_1
 - the 10.31.101.0/24 network mapped to the 10.11.101.0/24 network on FortiGate_2
- Configure an outgoing security policy with ordinary source NAT on both FortiGates.
- Configure an incoming security policy with the VIP as the destination on both FortiGates.
- Configure a route to the remote private network over the IPsec interface on both FortiGates.

To configure VIP mapping on both FortiGates

1. Go to **Policy & Objects > Virtual IPs** and create a new **Virtual IP**.
2. Enter the following information, and select **OK**:

Name	Enter a name, for example, <code>my_vip</code> .
External Interface	Select <code>FGT1_to_FGT2</code> . The IPsec interface.

VIP Type	Depending on both FortiGates, select one of the following options: <ul style="list-style-type: none"> • IPv4: If both FortiGates use IPv4 (Static NAT). • IPv6: If both FortiGates use IPv6 (Static NAT). • NAT46: Maps the IPv4 address into an IPv6 prefix. • NAT64: Maps the IPv6 address into an IPv4 prefix.
External IP Address/Range	For the External IP Address field enter: <p>10.21.101.1 when configuring FortiGate_1, or 10.31.101.1 when configuring FortiGate_2.</p>
Mapped IP Address/Range	For the Mapped IP Address enter 10.11.101.1. For the Range enter 10.11.101.254.
Port Forwarding	Disable

3. Repeat this procedure on both FortiGate_1 and FortiGate_2.

To configure the outbound security policy on both FortiGates

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information, and select **OK**:

Incoming Interface	Select Port 1 .
Outgoing Interface	Select FGT1_to_FGT2 . The IPsec interface.
Source	Select all .
Destination Address	Select all .
Action	Select ACCEPT
NAT	Enable NAT .

3. Repeat this procedure on both FortiGate_1 and FortiGate_2.

To configure the inbound security policy on both FortiGates

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information, and then select **OK**:

Incoming Interface	Select FGT1_to_FGT2 .
---------------------------	------------------------------

Outgoing Interface	Select Port 1 . The IPsec interface.
Source	Select all .
Destination Address	Select my-vip .
Action	Select ACCEPT
NAT	Disable NAT .

3. Repeat this procedure on both FortiGate_1 and FortiGate_2.

To configure the static route for both FortiGates

1. Go to **Network > Static Routes** and create a new **Route** (or **IPv6 Route** as necessary).
2. Enter the following information, and then select **OK**:

Destination	Enter a subnet of 10.31.101.0/24 when configuring FortiGate_1. Enter a subnet of 10.21.101.0/24 when configuring FortiGate_2.
Device	Select FGT1_to_FGT2 .
Gateway	Leave as default: 0.0.0.0.
Administrative Distance	Leave at default (10). If you have advanced routing on your network, you may have to change this value.
Advanced Options	If you have advanced routing on your network, enable Advanced Options and enter a Priority .

Solution for policy-based VPN

As with the route-based solution, users contact hosts at the other end of the VPN using an alternate subnet address. PC1 communicates with PC2 using IP address 10.31.101.10, and PC2 communicates with PC1 using IP address 10.21.101.10.

In this solution however, outbound NAT is used to translate the source address of packets from the 10.11.101.0/24 network to the alternate subnet address that hosts at the other end of the VPN use to reply. Inbound packets from the remote end have their destination addresses translated back to the 10.11.101.0/24 network.

For example, PC1 uses the destination address 10.31.101.10 to contact PC2. Outbound NAT on FortiGate_1 translates the PC1 source address to 10.21.101.10. At the FortiGate_2 end of the tunnel, the outbound NAT configuration translates the destination address to the actual PC2 address of 10.11.101.10. Similarly, PC2 replies to PC1 using destination address 10.21.101.10, with the PC2 source address translated to 10.31.101.10. PC1 and PC2 can communicate over the VPN even though they both have the same IP address.

You need to:

- Configure IPsec Phase 1 as you usually would for a policy-based VPN.
- Configure IPsec Phase 2 with the `use-natip disable` CLI option.
- Define a firewall address for the local private network, 10.11.101.0/24.
- Define a firewall address for the remote private network:
 - Define a firewall address for 10.31.101.0/24 on FortiGate_1
 - Define a firewall address for 10.21.101.0/24 on FortiGate_2
- Configure an outgoing IPsec security policy with outbound NAT to map 10.11.101.0/24 source addresses:
 - To the 10.21.101.0/24 network on FortiGate_1
 - To the 10.31.101.0/24 network on FortiGate_2

To configure IPsec Phase 2 - CLI

```
config vpn ipsec phase2
  edit "FGT1_FGT2_p2"
    set keepalive enable
    set pfs enable
    set phase1name FGT1_to_FGT2
    set proposal 3des-sha1 3des-md5
    set replay enable
    set use-natip disable
  end
```

In this example, your Phase 1 definition is named `FGT1_to_FGT2`. `use-natip` is set to `disable`, so you can specify the source selector using the **`src-addr-type`**, `src-start-ip / src-end-ip` or `src-subnet` keywords. This example leaves these keywords at their default values, which specify the subnet `0.0.0.0/0`.

The `pfs` keyword ensures that perfect forward secrecy (PFS) is used. This ensures that each Phase 2 key created is unrelated to any other keys in use.

To define the local private network firewall address

1. Go to **Policy & Objects > Addresses** and create a new **Address**.
2. Enter the following information and select **OK**.

Category	Set to Address .
Name	Enter <code>vpn-local</code> . A meaningful name for the local private network.
Type	Set to IP/Netmask .
Subnet / IP Range	10.11.101.0 255.255.255.0
Interface	Set to any .

To define the remote private network firewall address

1. Go to **Policy & Objects > Addresses** and create a new **Address**.
2. Enter the following information, and select **OK**:

Category	Set to Address .
----------	-------------------------

Name	Enter <code>vpn-remote</code> . A meaningful name for the remote private network.
Type	Set to IP/Netmask .
Subnet / IP Range	10.31.101.0 255.255.255.0 on FortiGate_1. 10.21.101.0 255.255.255.0 on FortiGate_2.
Interface	Any

To configure the IPsec security policy

In the CLI on FortiGate_1, enter the commands:

```
config firewall policy
  edit 1
    set srcintf "port1"
    set dstintf "port2"
    set srcaddr "vpn-local"
    set dstaddr "vpn-remote"
    set action ipsec
    set schedule "always"
    set service "ANY"
    set inbound enable
    set outbound enable
    set vpngateway "FGT1_to_FGT2"
    set natoutbound enable
    set natip 10.31.101.0 255.255.255.0
  end
```

Optionally, you can set everything except `natip` in the web-based manager and then use the CLI to set `natip`.

Enter the same commands on FortiGate_2, but set `natip` be 10.21.101.0 255.255.255.0.

Testing

The best testing is to look at the packets both as the VPN tunnel is negotiated, and when the tunnel is up.

Determining what the other end of the VPN tunnel is proposing

1. Start a terminal program such as PuTTY and set it to log all output.
When necessary refer to the logs to locate information when output is verbose.
2. Logon to the FortiGate unit using a `super_admin` account.
3. Enter the following CLI commands.
4. Display all the possible IKE error types and the number of times they have occurred:

```
diag vpn ike errors
```

5. Check for existing debug sessions:

```
diag debug info
```

If a debug session is running, to halt it enter:

```
diag debug disable
```

6. Confirm your proposal settings:

```
diag vpn ike config list
```

7. If your proposal settings do not match what you expect, make a change to it and save it to force an update in memory. If that fixes the problem, stop here.

8. List the current vpn filter:

```
diag vpn ike filter
```

9. If all fields are set to any, there are no filters set and all VPN IKE packets will be displayed in the debug output. If your system has only a few VPNs, skip setting the filter.
If your system has many VPN connections this will result in very verbose output and make it very difficult to locate the correct connection attempt.

10. Set the VPN filter to display only information from the destination IP address for example 10.10.10.10:

```
diag vpn ike log-filter dst-addr4 10.10.10.10
```

To add more filter options, enter them one per line as above. Other filter options are:

clear	erase the current filter
dst-addr6	the IPv6 destination address range to filter by
dst-port	the destination port range to filter by
interface	interface that IKE connection is negotiated over
list	display the current filter
name	the phase1 name to filter by
negate	negate the specified filter parameter
src-addr4	the IPv4 source address range to filter by
src-addr6	the IPv6 source address range to filter by
src-port	the source port range to filter by
vd	index of virtual domain. 0 matches all

11. Start debugging:

```
diag debug app ike 255  
diag debug enable
```

12. Have the remote end attempt a VPN connection.

If the remote end attempts the connection they become the initiator. This situation makes it easier to debug VPN tunnels because then you have the remote information and all of your local information. by initiate the connection, you will not see the other end's information.

13. If possible go to the web-based manager on your FortiGate unit, go to the VPN monitor and try to bring the tunnel up.
14. Stop the debug output:

```
diag debug disable
```

15. Go back through the output to determine what proposal information the initiator is using, and how it is different from your VPN P1 proposal settings.

Things to look for in the debug output of attempted VPN connections are shown below.

Important terms to look for in VPN debug output

initiator	Starts the VPN attempt, in the above procedure that is the remote end
responder	Answers the initiator's request
local ID	In aggressive mode, this is not encrypted
error no SA proposal chosen	There was no proposal match — there was no encryption-authentication pair in common, usually occurs after a long list of proposal attempts
R U THERE and R U THERE ack	dead peer detection (dpd), also known as dead gateway detection — after three failed attempts to contact the remote end it will be declared dead, no farther attempts will be made to contact it
negotiation result	lists the proposal settings that were agreed on
SA_life_soft and SA_life_hard	negotiating a new key, and the key life
R U THERE	If you see this, it means Phase 1 was successful
tunnel up	the negotiation was successful, the VPN tunnel is operational

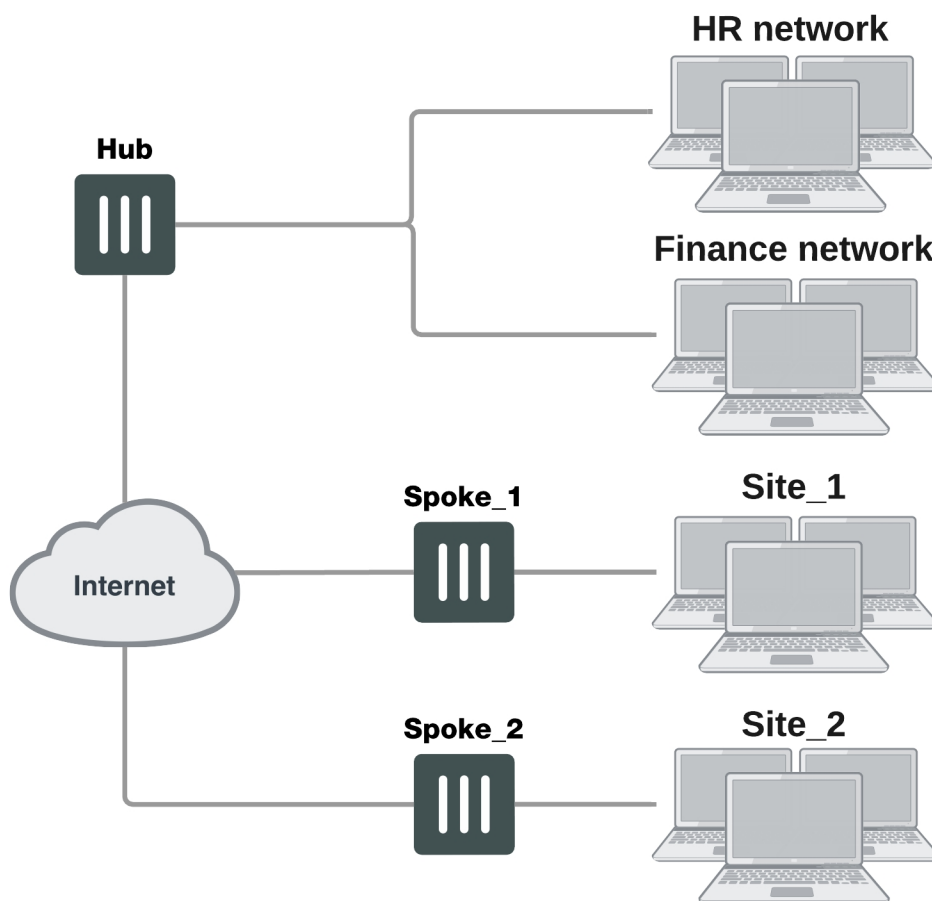
Hub-and-spoke configurations

This section describes how to set up hub-and-spoke IPsec VPNs. The following topics are included in this section:

Configuration overview

In a hub-and-spoke configuration, VPN connections radiate from a central FortiGate unit (the hub) to a number of remote peers (the spokes). Traffic can pass between private networks behind the hub and private networks behind the remote peers. Traffic can also pass between remote peer private networks through the hub.

Example hub-and-spoke configuration



The actual implementation varies in complexity depending on:

- Whether the spokes are statically or dynamically addressed
- The addressing scheme of the protected subnets
- How peers are authenticated

This guide discusses the issues involved in configuring a hub-and-spoke VPN and provides some basic configuration examples.

Hub-and-spoke infrastructure requirements

- The FortiGate hub must be operating in NAT mode and have a static public IP address.
- Spokes may have static IP addresses, dynamic IP addresses (see [FortiGate dialup-client configurations on page 1745](#)), or static domain names and dynamic IP addresses (see [Dynamic DNS configuration on page 1723](#)).

Spoke gateway addressing

The public IP address of the spoke is the VPN remote gateway as seen from the hub. Statically addressed spokes each require a separate VPN Phase 1 configuration on the hub. When there are many spokes, this becomes rather cumbersome.

Using dynamic addressing for spokes simplifies the VPN configuration because then the hub requires only a single Phase 1 configuration with “dialup user” as the remote gateway. You can use this configuration even if the remote peers have static IP addresses. A remote peer can establish a VPN connection regardless of its IP address if its traffic selectors match and it can authenticate to the hub. See [Configuration overview on page 1702](#) for an example of this configuration.

Protected networks addressing


The addresses of the protected networks are needed to configure destination selectors and sometimes for security policies and static routes. The larger the number of spokes, the more addresses there are to manage. You can

- Assign spoke subnets as part of a larger subnet, usually on a new network
- or
- Create address groups that contain all of the needed addresses

Using aggregated subnets

If you are creating a new network, where subnet IP addresses are not already assigned, you can simplify the VPN configuration by assigning spoke subnets that are part of a large subnet.

Aggregated subnets



All spokes use the large subnet address, 10.1.0.0/16 for example, as:

- The IPsec destination selector
- The destination of the security policy from the private subnet to the VPN (required for policy-based VPN, optional for route-based VPN)
- The destination of the static route to the VPN (route-based)

Each spoke uses the address of its own protected subnet as the IPsec source selector and as the source address in its VPN security policy. The remote gateway is the public IP address of the hub FortiGate unit.

Using an address group

If you want to create a hub-and-spoke VPN between existing private networks, the subnet addressing usually does not fit the aggregated subnet model discussed earlier. All of the spokes and the hub will need to include the addresses of all the protected networks in their configuration.

On FortiGate units, you can define a named firewall address for each of the remote protected networks and add these addresses to a firewall address group. For a policy-based VPN, you can then use this address group as the destination of the VPN security policy.

For a route-based VPN, the destination of the VPN security policy can be set to All. You need to specify appropriate routes for each of the remote subnets.

Authentication

Authentication is by a common pre-shared key or by certificates. For simplicity, the examples in this chapter assume that all spokes use the same pre-shared key.

Configure the hub

At the FortiGate unit that acts as the hub, you need to:

- Configure the VPN to each spoke
- Configure communication between spokes

You configure communication between spokes differently for a policy-based VPN than for a route-based VPN. For a policy-based VPN, you configure a VPN concentrator. For a route-based VPN, you must either define security policies or group the IPsec interfaces into a zone.

Define the hub-spoke VPNs

Perform these steps at the FortiGate unit that will act as the hub. Although this procedure assumes that the spokes are all FortiGate units, a spoke could also be VPN client software, such as FortiClient Endpoint Security.

Configuring the VPN hub

1. At the hub, define the Phase 1 configuration for each spoke. See [Phase 1 parameters on page 1655](#). Enter these settings in particular:

Name	Enter a name to identify the VPN in Phase 2 configurations, security policies and the VPN monitor.
Remote Gateway	<p>The remote gateway is the other end of the VPN tunnel. There are three options:</p> <p>Static IP Address — Enter the spoke's public IP Address. You will need to create a Phase 1 configuration for each spoke. Either the hub or the spoke can establish the VPN connection.</p> <p>Dialup User — No additional information is needed. The hub accepts connections from peers with appropriate encryption and authentication settings. Only one Phase 1 configuration is needed for multiple dialup spokes. Only the spoke can establish the VPN tunnel.</p> <p>Dynamic DNS — If the spoke subscribes to a dynamic DNS service, enter the spoke's Dynamic DNS domain name. Either the hub or the spoke can establish the VPN connection. For more information, see Dynamic DNS configuration on page 1.</p>
Local Interface	Select the FortiGate interface that connects to the remote gateway. This is usually the FortiGate unit's public interface.

2. Define the Phase 2 parameters needed to create a VPN tunnel with each spoke. See [Phase 2 parameters on page 1675](#). Enter these settings in particular:

Name	Enter a name to identify this spoke Phase 2 configuration.
Phase 1	Select the name of the Phase 1 configuration that you defined for this spoke.

IPsec VPN in ADVPN hub-and-spoke

IPsec VPN traffic is allowed through a tunnel between an ADVPN hub-and-spoke.

CLI syntax:

```
config vpn ipsec phase1-interface
edit "int-fgtb"
```



```

...
set auto-discovery-sender [enable | disable]
set auto-discovery-receiver [enable | disable]
set auto-discovery-forwarder [enable | disable]
...
next
end
config vpn ipsec phase2-interface
edit "int-fgtb"
...
set auto-discovery-sender phase1 [enable | disable]
...
next
end

```

Define the hub-spoke security policies

1. Define a name for the address of the private network behind the hub. For more information, see [Defining policy addresses on page 1](#).
2. Define names for the addresses or address ranges of the private networks behind the spokes. For more information, see [Defining policy addresses on page 1](#).
3. Define the VPN concentrator. See [To define the VPN concentrator on page 1707](#).
4. Define security policies to permit communication between the hub and the spokes. For more information, see [Defining VPN security policies on page 1](#).

Route-based VPN security policies

Define ACCEPT security policies to permit communications between the hub and the spoke. You need one policy for each direction.

Adding policies

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Leave the **Policy Type** as **Firewall** and leave the **Policy Subtype** as **Address**.
3. Enter these settings in particular:

Incoming Interface	Select the VPN Tunnel (IPsec Interface) you configured in Step 1.
Source Address	Select the address name you defined in Step 2 for the private network behind the spoke FortiGate unit.
Outgoing Interface	Select the hub's interface to the internal (private) network.
Destination Address	Select the source address that you defined in Step 1.
Action	Select ACCEPT .
Enable NAT	Enable.

Incoming Interface	Select the VPN Tunnel (IPsec Interface) you configured in Step 1.
---------------------------	---

Source Address	Select the address name you defined in Step 2 for the private network behind the spoke FortiGate units.
Outgoing Interface	Select the source address that you defined in Step 1.
Destination Address	Select the hub's interface to the internal (private) network.
Action	Select ACCEPT .
Enable NAT	Enable.

Policy-based VPN security policy

Define an IPsec security policy to permit communications between the hub and the spoke.

Adding policies

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter these settings in particular:

Incoming Interface	Select the hub's interface to the internal (private) network.
Source Address	Select the source address that you defined in Step 1.
Outgoing Interface	Select the hub's public network interface.
Destination Address	Select the address name you defined in Step 2 for the private network behind the spoke FortiGate unit.
VPN Tunnel	Select Use Existing and select the name of the Phase 1 configuration that you created for the spoke in Step 1. Select Allow traffic to be initiated from the remote site to enable traffic from the remote network to initiate the tunnel.

In the policy list, arrange the policies in the following order:

- IPsec policies that control traffic between the hub and the spokes first
- The default security policy last

Configuring communication between spokes (policy-based VPN)

For a policy-based hub-and-spoke VPN, you define a concentrator to enable communication between the spokes.

To define the VPN concentrator

1. At the hub, go to **VPN > IPsec Concentrator** and select **Create New**.
2. In the **Concentrator Name** field, type a name to identify the concentrator.
3. From the **Available Tunnels** list, select a VPN tunnel and then select the right-pointing arrow.
4. Repeat Step 3 until all of the tunnels associated with the spokes are included in the concentrator.
5. Select **OK**.

Configuring communication between spokes (route-based VPN)

For a route-based hub-and-spoke VPN, there are several ways you can enable communication between the spokes:

- Put all of the IPsec interfaces into a zone and enable intra-zone traffic. This eliminates the need for any security policy for the VPN, but you cannot apply UTM features to scan the traffic for security threats.
- Put all of the IPsec interfaces into a zone and create a single zone-to-zone security policy
- Create a security policy for each pair of spokes that are allowed to communicate with each other. The number of policies required increases rapidly as the number of spokes increases.

Using a zone as a concentrator

A simple way to provide communication among all of the spokes is to create a zone and allow intra-zone communication. You cannot apply UTM features using this method.

1. Go to **Network > Interfaces**.
2. Select the down-arrow on the **Create New** button and select **Zone**.
3. In the **Zone Name** field, enter a name, such as `Our_VPN_zone`.
4. Clear **Block intra-zone traffic**.
5. In the **Interface Members** list, select the IPsec interfaces that are part of your VPN.
6. Select **OK**.

Using a zone with a policy as a concentrator

If you put all of the hub IPsec interfaces involved in the VPN into a zone, you can enable communication among all of the spokes and apply UTM features with just one security policy.

Creating a zone for the VPN

1. Go to **Network > Interfaces**.
2. Select the down-arrow on the **Create New** button and select **Zone**.
3. In the **Zone Name** field, enter a name, such as `Our_VPN_zone`.
4. Select **Block intra-zone traffic**.
5. In the **Interface Members** list, select the IPsec interfaces that are part of your VPN.
6. Select **OK**.

Creating a security policy for the zone

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Leave the **Policy Type** as **Firewall** and leave the **Policy Subtype** as **Address**.
3. Enter the settings: and select **OK**.

Incoming Interface	Select the zone you created for your VPN.
Source Address	Select All .
Outgoing Interface	Select the zone you created for your VPN.

Destination Address	Select All .
Action	Select ACCEPT .
Enable NAT	Enable.

Using security policies as a concentrator

To enable communication between two spokes, you need to define an ACCEPT security policy for them. To allow either spoke to initiate communication, you must create a policy for each direction. This procedure describes a security policy for communication from Spoke 1 to Spoke 2. Others are similar.

1. Define names for the addresses or address ranges of the private networks behind each spoke. For more information, see [Defining policy addresses on page 1](#).
2. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
3. Leave the **Policy Type** as **Firewall** and leave the **Policy Subtype** as **Address**.
4. Enter the settings and select **OK**.

Incoming Interface	Select the IPsec interface that connects to Spoke 1.
Source Address	Select the address of the private network behind Spoke 1.
Outgoing Interface	Select the IPsec interface that connects to Spoke 2.
Destination Address	Select the address of the private network behind Spoke 2.
Action	Select ACCEPT .
Enable NAT	Enable.

Configure the spokes

Although this procedure assumes that the spokes are all FortiGate units, a spoke could also be VPN client software, such as FortiClient Endpoint Security.

Perform these steps at each FortiGate unit that will act as a spoke.

Creating the Phase 1 and phase_2 configurations

1. At the spoke, define the Phase 1 parameters that the spoke will use to establish a secure connection with the hub. See [Phase 1 parameters on page 1655](#). Enter these settings:

Remote Gateway	Select Static IP Address .
IP Address	Type the IP address of the interface that connects to the hub.

2. Create the Phase 2 tunnel definition. See [Phase 2 parameters on page 1675](#). Select the set of Phase 1 parameters that you defined for the hub. You can select the name of the hub from the **Static IP Address** part of the list.

Configuring security policies for hub-to-spoke communication

1. Create an address for this spoke. See [Defining policy addresses on page 1](#). Enter the IP address and netmask of the private network behind the spoke.
2. Create an address to represent the hub. See [Defining policy addresses on page 1](#). Enter the IP address and netmask of the private network behind the hub.
3. Define the security policy to enable communication with the hub.

Route-based VPN security policy

Define two security policies to permit communications to and from the hub.

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Leave the **Policy Type** as **Firewall** and leave the **Policy Subtype** as **Address**.
3. Enter these settings:

Incoming Interface	Select the virtual IPsec interface you created.
Source Address	Select the hub address you defined in Step 1.
Outgoing Interface	Select the spoke's interface to the internal (private) network.
Destination Address	Select the spoke addresses you defined in Step 2.
Action	Select ACCEPT .
Enable NAT	Enable

Incoming Interface	Select the spoke's interface to the internal (private) network.
Source Address	Select the spoke address you defined in Step 1.
Outgoing Interface	Select the virtual IPsec interface you created.
Destination Address	Select the hub destination addresses you defined in Step 2.
Action	Select ACCEPT .
Enable NAT	Enable

Policy-based VPN security policy

Define an IPsec security policy to permit communications with the hub. See [Defining VPN security policies on page 1](#).

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter these settings in particular:

Incoming Interface	Select the spoke's interface to the internal (private) network.
Source Address	Select the spoke address you defined in Step 1.

Outgoing Interface	Select the spoke's interface to the external (public) network.
Destination Address	Select the hub address you defined in Step 2.
VPN Tunnel	Select Use Existing and select the name of the Phase 1 configuration you defined. Select Allow traffic to be initiated from the remote site to enable traffic from the remote network to initiate the tunnel.

Configuring security policies for spoke-to-spoke communication

Each spoke requires security policies to enable communication with the other spokes. Instead of creating separate security policies for each spoke, you can create an address group that contains the addresses of the networks behind the other spokes. The security policy then applies to all of the spokes in the group.

1. Define destination addresses to represent the networks behind each of the other spokes. Add these addresses to an address group.
2. Define the security policy to enable communication between this spoke and the spokes in the address group you created.

Policy-based VPN security policy

Define an IPsec security policy to permit communications with the other spokes. See [Defining VPN security policies on page 1](#). Enter these settings in particular:

Route-based VPN security policy

Define two security policies to permit communications to and from the other spokes.

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Leave the **Policy Type** as **Firewall** and leave the **Policy Subtype** as **Address**.
3. Enter these settings in particular:

Incoming Interface	Select the virtual IPsec interface you created.
Source Address	Select the spoke address group you defined in Step " Configure the spokes " on page 1709.
Outgoing Interface	Select the spoke's interface to the internal (private) network.
Destination Address	Select this spoke's address name.
Action	Select ACCEPT .
Enable NAT	Enable

4. Select **Create New**, leave the **Policy Type** as **Firewall** and leave the **Policy Subtype** as **Address**, and enter these settings:

Incoming Interface	Select the spoke's interface to the internal (private) network.
---------------------------	---

Source Address	Select this spoke's address name.
Outgoing Interface	Select the virtual IPsec interface you created.
Destination Address	Select the spoke address group you defined in Step 1.
Action	Select ACCEPT .
Enable NAT	Enable

Policy-based VPN security policy

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following:

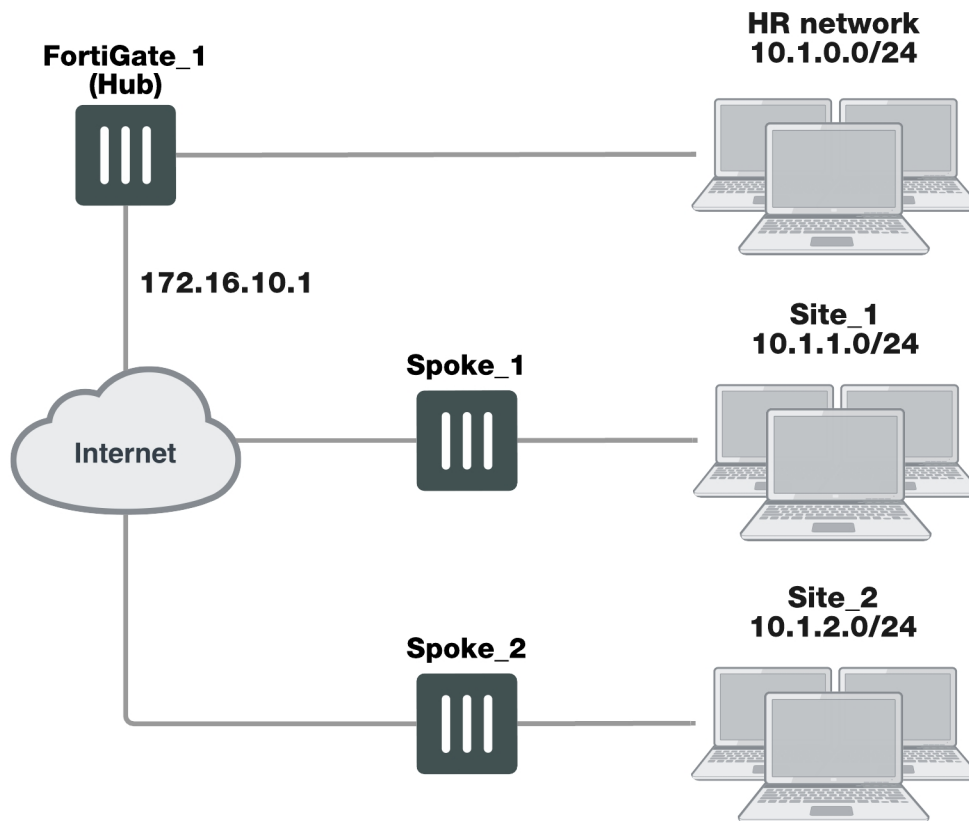
Incoming Interface	Select this spoke's internal (private) network interface.
Source Address	Select this spoke's source address.
Outgoing Interface	Select the spoke's interface to the external (public) network.
Destination Address	Select the spoke address group you defined in Step 1.
VPN Tunnel	<p>Select Use Existing and select the name of the Phase 1 configuration you defined.</p> <p>Select Allow traffic to be initiated from the remote site to enable traffic from the remote network to initiate the tunnel.</p>

Place this policy or policies in the policy list above any other policies having similar source and destination addresses.

Dynamic spokes configuration example

This example demonstrates how to set up a basic route-based hub-and-spoke IPsec VPN that uses preshared keys to authenticate VPN peers.

Example hub-and-spoke configuration



In the example configuration, the protected networks 10.1.0.0/24, 10.1.1.0/24 and 10.1.2.0/24 are all part of the larger subnet 10.1.0.0/16. The steps for setting up the example hub-and-spoke configuration create a VPN among Site 1, Site 2, and the HR Network.

The spokes are dialup. Their addresses are not part of the configuration on the hub, so only one spoke definition is required no matter the number of spokes. For simplicity, only two spokes are shown.

In an ADVPN topology, any two pair of peers can create a shortcut, as long as one of the devices is not behind NAT.

The on-the-wire format of the ADVPN messages use TLV encoding. Because of this, this feature is not compatible with any previous ADVPN builds.

Configure the hub (FortiGate_1)

The Phase 1 configuration defines the parameters that FortiGate_1 will use to authenticate spokes and establish secure connections.

For the purposes of this example, one preshared key will be used to authenticate all of the spokes. Each key must contain at least 6 printable characters and best practices dictates that it only be known by network administrators. For optimum protection against currently known attacks, each key must consist of a minimum of 16 randomly chosen alphanumeric characters.

Define the IPsec configuration

1. At FortiGate_1, go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
2. Edit the **Phase 1 Proposal** (if it is not available, you may need to click the **Convert to Custom Tunnel** button). Define the Phase 1 parameters that the hub will use to establish a secure connection to the spokes.

Name	Enter a name (for example, <code>toSpokes</code>).
Remote Gateway	Dialup user
Local Interface	External
Mode	Main
Authentication Method	Preshared Key
Pre-shared Key	Enter the preshared key.
Peer Options	Any peer ID

The basic Phase 2 settings associate IPsec Phase 2 parameters with the Phase 1 configuration and specify the remote end points of the VPN tunnels.

3. Open the **Phase 2 Selectors** panel (if it is not available, you may need to click the **Convert to Custom Tunnel** button).
4. Enter the following information, and select **OK**:

Name	Enter a name for the Phase 2 definition (for example, <code>toSpokes_ph2</code>).
Phase 1	Select the Phase 1 configuration that you defined previously (for example, <code>toSpokes</code>).

Define the security policies

security policies control all IP traffic passing between a source address and a destination address. For a route-based VPN, the policies are simpler than for a policy-based VPN. Instead of an IPSEC policy, you use an ACCEPT policy with the virtual IPsec interface as the external interface.

Before you define security policies, you must first define firewall addresses to use in those policies. You need addresses for:

- The HR network behind FortiGate_1
- The aggregate subnet address for the protected networks

Defining the IP address of the HR network behind FortiGate_1

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**, enter the following information, and select **OK**:

Name	Enter an address name (for example, <code>HR_Network</code>).
-------------	--

Type	Subnet
Subnet/IP Range	Enter the IP address of the HR network behind FortiGate_1 (for example, 10.1.0.0/24).

Specifying the IP address the aggregate protected subnet

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**, enter the following information, and select **OK**:

Address Name	Enter an address name (for example, Spoke_net).
Type	Subnet
Subnet/IP Range	Enter the IP address of the aggregate protected network, 10.1.0.0/16

Defining the security policy for traffic from the hub to the spokes

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**,
2. Leave the **Policy Type** as **Firewall** and leave the **Policy Subtype** as **Address**.
3. Enter the following information, and select **OK**:

Incoming Interface	Select the interface to the HR network, port 1 .
Source Address	Select HR_Network .
Outgoing Interface	Select the virtual IPsec interface that connects to the spokes, toSpokes .
Destination Address	Select Spoke_net .
Action	Select ACCEPT .

Place the policy in the policy list above any other policies having similar source and destination addresses.

Configure communication between spokes

Spokes communicate with each other through the hub. You need to configure the hub to allow this communication. An easy way to do this is to create a zone containing the virtual IPsec interfaces even if there is only one, and create a zone-to-zone security policy.

1. Go to **Network > Interfaces**.
2. Select the down-arrow on the **Create New** button and select **Zone**.
3. In the **Zone Name** field, enter a name, such as **Our_VPN_zone**.
4. Select **Block intra-zone traffic**.
You could enable intra-zone traffic and then you would not need to create a security policy. But, you would not be able to apply UTM features.
5. In **Interface Members**, select the virtual IPsec interface, **toSpokes**.
6. Select **OK**.

Creating a security policy for the zone

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Leave the **Policy Type** as **Firewall** and leave the **Policy Subtype** as **Address**.
3. Enter these settings:

Incoming Interface	Select <code>Our_VPN_zone</code> .
Source Address	Select All .
Outgoing Interface	Select <code>Our_VPN_zone</code> .
Destination Address	Select All .
Action	Select ACCEPT .
Enable NAT	Enable.

4. Select **OK**.

Configure the spokes

In this example, all spokes have nearly identical configuration, requiring the following:

- Phase 1 authentication parameters to initiate a connection with the hub.
- Phase 2 tunnel creation parameters to establish a VPN tunnel with the hub.
- A source address that represents the network behind the spoke. This is the only part of the configuration that is different for each spoke.
- A destination address that represents the aggregate protected network.
- A security policy to enable communications between the spoke and the aggregate protected network

Define the IPsec configuration

At each spoke, create the following configuration.

1. At the spoke, go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
2. Edit the **Phase 1 Proposal** (if it is not available, you may need to click the **Convert to Custom Tunnel** button). Enter the following information:

Name	Type a name, for example, <code>toHub</code> .
Remote Gateway	Select Static IP Address .
IP Address	Enter <code>172.16.10.1</code> .
Local Interface	Select Port2 .
Mode	Main
Authentication Method	Preshared Key
Pre-shared Key	Enter the preshared key. The value must be identical to the preshared key that you specified previously in the <code>FortiGate_1</code> configuration

Peer Options	Select Any peer ID .
---------------------	-----------------------------

1. Open the **Phase 2 Selectors** panel (if it is not available, you may need to click the **Convert to Custom Tunnel** button).
2. Enter the following information and select **OK**:

Name	Enter a name for the tunnel, for example, <code>toHub_ph2</code> .
Phase 1	Select the name of the Phase 1 configuration that you defined previously, for example, <code>toHub</code> .
Advanced	Select to show the following Quick Mode Selector settings.
Source	Enter the address of the protected network at this spoke. For <code>spoke_1</code> , this is <code>10.1.1.0/24</code> . For <code>spoke_2</code> , this is <code>10.1.2.0/24</code> .
Destination	Enter the aggregate protected subnet address, <code>10.1.0.0/16</code> .

Define the security policies

You need to define firewall addresses for the spokes and the aggregate protected network and then create a security policy to enable communication between them.

Defining the IP address of the network behind the spoke

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New** and enter the following information:

Address Name	Enter an address name, for example <code>LocalNet</code> .
Type	Subnet
Subnet/IP Range	Enter the IP address of the private network behind the spoke. For <code>spoke_1</code> , this is <code>10.1.1.0/24</code> . For <code>spoke_2</code> , this is <code>10.1.2.0/24</code> .

Specifying the IP address of the aggregate protected network

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New** and enter the following information:

Address Name	Enter an address name, for example, <code>Spoke_net</code> .
Type	Subnet
Subnet/IP Range	Enter the IP address of the aggregate protected network, <code>10.1.0.0/16</code> .

Defining the security policy

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Leave the **Policy Type** as **Firewall** and leave the **Policy Subtype** as **Address**.
3. Enter the following information:

Incoming Interface	Select the virtual IPsec interface, <code>toHub</code> .
Source Address	Select the aggregate protected network address <code>Spoke_net</code> .
Outgoing Interface	Select the interface to the internal (private) network, <code>port1</code> .
Destination Address	Select the address for this spoke's protected network <code>LocalNet</code> .
Action	Select ACCEPT .

4. Select **Create New**.
5. Leave the **Policy Type** as **Firewall** and leave the **Policy Subtype** as **Address**.
6. Enter the following information, and select **OK**:

Incoming Interface	Select the interface to the internal private network, <code>port1</code> .
Source Address	Select the address for this spoke's protected network, <code>LocalNet</code> .
Outgoing Interface	Select the virtual IPsec interface, <code>toHub</code> .
Destination Address	Select the aggregate protected network address, <code>Spoke_net</code> .
Action	Select ACCEPT .

Place these policies in the policy list above any other policies having similar source and destination addresses.

One-Click VPN (OCVPN)

One-Click VPN (OCVPN) is a cloud-based solution that greatly simplifies the provisioning and configuration of IPsec VPN. The administrator enables OCVPN with a single click, adds the required subnets, and then the configuration is complete. The OCVPN updates each FortiGate automatically as devices join/leave the VPN, as subnets are added/removed, when dynamic external IPs change (e.g. DHCP/PPPoE), and when WAN interface bindings change (as in the dual WAN redundancy case).

Configuration changes and events are automatically propagated across participating nodes without user intervention, so in a sense, the VPN manages itself as a unit with only bare minimum user input. The user specifies which subnets to participate in the VPN. Everything else happens transparently to the user.

After registering devices with FortiCare, devices use SSL to register local subnets with the OCVPN cloud service at <https://productapi.fortinet.com>. The WAN IP is determined automatically (devices must use a publicly routed external WAN IP address) and the gateway IP address and participating subnets are uploaded to a cloud repository that collects and stores the information in each customer's FortiCare account.

The following limitations apply to FortiOS OCVPN:

- The FortiGate must be registered with a valid FortiCare Support license.
- Only full-mesh VPN configurations using PSK cryptography are supported.
- Public IPs must be used (FortiGates behind NAT cannot participate).
- Non-root VDOMs and FortiGate VMs are not supported.
- Up to 16 nodes can be added to the OCVPN cloud, each with a maximum of 16 subnets.

OCVPN support for High Availability (HA)

As of 6.0.2, HA-enabled devices are now supported by OCVPN.

Prior to establishing the HA cluster, if OCVPN is in use then both devices should be registered to the OCVPN cloud service. During failover, the old serial number is withdrawn and a new serial number (and VPN) is added, to account for the change in status.

General configuration

If FortiCare Support is registered on the FortiGate, you can configure OCVPN in FortiOS under **VPN > One-Click VPN Settings**.

Once enabled, you can add the relevant **Subnets**, as well as view any **Cloud Members** currently participating in the cloud-serviced VPN (you may need to **Refresh** the **Cloud Members** table).

If you wish to change the polling interval, you must do so in the CLI Console (see below).

To enable and configure OCVPN - CLI:

```
config vpn ocvpn
  set status {enable | disable}
  set poll interval <30 - 120>
  config subnets
    edit 1
      set subnet 10.1.1.0 255.255.255.0
```

```

    next
    edit 2
        set subnet 10.1.2.0 255.255.255.0
    next
end
end

```

where:

Command	Description
<code>set status {enable disable}</code>	This command enables or disables the service. After a device has been registered with FortiCare, enabling this feature registers the device with the OCVPN cloud service. Disabling causes the device to be unregistered, and removed from the table of VPN members.
<code>poll interval <30 - 120></code>	Set the OCVPN polling interval. Enter an integer value from 30 to 120 (default = 60).
<code>config subnets</code>	This is the OCVPN subcommand for configuring the list of participating subnets.

Key exchange

Keys are generated automatically by OCVPN, but without explicit acknowledgment and state management, it would be impossible for the cloud to destroy keys after distribution to customer devices. Permanently storing customer keys in the cloud is undesirable for a host of reasons, so the RegAck request was introduced to effectively address the problem and allow the cloud to destroy keys after they have been installed. One key is generated per customer. When a new member joins, a new key is generated and distributed to all group members at the next poll interval (the default is 60 seconds).

Authentication is handled by SSL and proof of identity is established by the device serial number in the signed RSA certificate. The SN is sent in all messages to the cloud.

If you have a FortiWeb server performing authentication, the process is different. Since the OCVPN microservice doesn't run on the FortiWeb server, OCVPN authentication and secure segregation of customer data is handled as follows:

- FortiWeb extracts the ASN1 CN from the certificate and attaches it to the decrypted HTTP messages forwarded to OCVPN.
- OCVPN checks the presented device SN against the SN included in the certificate ID.
- If they don't match, OCVPN returns '401 Unauthorized' and the authentication transaction is cancelled.

Device polling and controller information

Instead of a central controller actively directing and pushing out the devices in response to network topology changes, FortiOS architecture uses device polling to propagate changes across nodes in the VPN. State changes are tracked carefully across the system so all devices always have the same view of the network (with some delay in propagating changes due to polling). Similarly the OCVPN cloud always know the state of each device. This is

essential to being able to manage the keys properly, and be able to discard them after they have been installed on each device.

The control layer is implemented on each device as a state machine, where information is translated from the member table into a working configuration—with IPsec phase1 and phase2 objects with default parameters, firewall address and address group objects, firewall policies, and static routes. The resulting configuration may be edited normally, e.g. DPD settings, DH group, crypto transform, firewall policy profiles for AV/IPS, etc. This is to provide a level of flexibility and usability.

The control layer's responsibility is to ensure that the network data on any device, and by extension the configuration, always stays in sync with the network view stored in the cloud, and in sync with all the other devices, regardless of intermittent network errors that could occur at any point in the system. The system is designed to handle network errors, changes, and events and keep the IPsec configuration consistently and reliable in sync.

Configuration information is managed in a fixed table: 16 nodes maximum, 16 subnets per node. After the table is populated, full mesh configuration is calculated and installed into the CMDB.

System states

The system is stateless across reboots. It re-registers after reboot, which re-initializes the state of the system. After bootup, the system is stateful across changes and polling interval queries/updates. The state file contains the hostname, current WAN ifname, current WAN IP, assigned slot, current state, previous state, current OCVPN table revision, last OCVPN response code (register/update), last polling response code, number of members, current member bitmask, previous member bitmask. The system uses this state information to track state changes locally and in the cloud.

Possible device states are:

```
enum cvpn_state {
    cvpn_st_none,
    cvpn_st_unregistered,
    cvpn_st_registering,
    cvpn_st_updating,
    cvpn_st_unregistering,
    cvpn_st_acknowledging,
    cvpn_st_registered
};
```

A normal sequence would be registering (updating) -> acknowledging -> registered.

Even though SSL/TCP is stateful and ensures delivery, the OCVPN microservice doesn't run on a FortiWeb SSL termination server. See ["Key exchange" on page 1720](#) for more info about how FortiWeb configuration differs. The explicit acknowledgment message (RegAck) ensures the OCVPN service knows when all nodes have received and applied the latest revision of the network information and key.

Debugging and logging

OCVPN debugging and logging is handled through a common API function. All debugs (except polling) are logged to /tmp/ocvpn/log. When the size of the log file exceeds 128k, the file is truncated and only the most recent 32k is saved.

The following diagnose commands may be useful when troubleshooting and debugging OCVPN configurations.

Command	Description
<code>diag vpn ocvpn</code>	Top level diagnose command for OCVPN.
<code>device-state</code>	Display OCVPN device state.
<code>log</code>	Display OCVPN log file from the device.
<code>status</code>	Display the current status of the device and last response code from the OCVPN service.
<code>print-members</code>	Print the OCVPN member table. This command accesses the OCVPN cloud service to retrieve the latest information, irrespective of the state of the device. It prints the raw JSON responses from OCVPN.

Dynamic DNS configuration

This section describes how to configure a site-to-site VPN, in which one FortiGate unit has a static IP address and the other FortiGate unit has a domain name and a dynamic IP address.

The following topics are included in this section:

Dynamic DNS over VPN concepts

A typical computer has a static IP address and one or more DNS servers to resolve fully qualified domain names (FQDN) into IP addresses. A domain name assigned to this computer is resolved by any DNS server having an entry for the domain name and its static IP address. The IP address never changes or changes only rarely so the DNS server can reliably say it has the correct address for that domain all the time.

Dynamic DNS (DDNS)

It is different when a computer has a dynamic IP address, such as an IP address assigned dynamically by a DHCP server, and a domain name. Computers that want to contact this computer do not know what its current IP address is. To solve this problem there are dynamic DNS (DDNS) servers. These are public servers that store a DNS entry for your computer that includes its current IP address and associated domain name. These entries are kept up to date by your computer sending its current IP address to the DDNS server to ensure its entry is always up to date. When other computers want to contact your domain, their DNS gets your IP address from your DDNS server. To use DDNS servers, you must subscribe to them and usually pay for their services.

When configuring DDNS on your FortiGate unit, go to **Network > DNS** and enable **Enable FortiGuard DDNS**. Then select the interface with the dynamic connection, which DDNS server you have an account with, your domain name, and account information. If your DDNS server is not on the list, there is a generic option where you can provide your DDNS server information.

Routing

When an interface has some form of changing IP address (DDNS, PPPoE, or DHCP assigned address), routing needs special attention. The standard static route cannot handle the changing IP address. The solution is to use the dynamic-gateway command in the CLI. Say for example you already have four static routes, and you have a PPPoE connection over the wan2 interface and you want to use that as your default route.

The route is configured on the dynamic address VPN peer trying to access the static address FortiGate unit.

Configuring dynamic gateway routing - CLI

```
config router static
  edit 5
    set dst 0.0.0.0 0.0.0.0
    set dynamic-gateway enable
    set device wan2
  next
end
```

For more information on DDNS, see the [System Administration](#) handbook chapter.

DDNS over VPN

IPsec VPN expects an IP address for each end of the VPN tunnel. All configuration and communication with that tunnel depends on the IP addresses as reference points. However, when the interface the tunnel is on has DDNS enabled there is no set IP address. The remote end of the VPN tunnel now needs another way to reference your end of the VPN tunnel. This is accomplished using Local ID.

A FortiGate unit that has a domain name and a dynamic IP address can initiate VPN connections anytime. The remote peer can reply to the local FortiGate unit using the source IP address that was sent in the packet header because it is current. Without doing a DNS lookup first, the remote peer runs the risk of the dynamic IP changing before it attempts to connect. To avoid this, the remote peer must perform a DNS lookup for the domain name of to be sure of the dynamic IP address before initiating the connection.

Remote gateway

When configuring the Phase 1 entry for a VPN tunnel, the Remote Gateway determines the addressing method the remote end of the tunnel uses as one of Static IP Address, Dialup User, or Dynamic DNS. There are different fields for each option.

When you select the Dynamic DNS VPN type there is a related field called Dynamic DNS. The Dynamic DNS field is asking for the FQDN of the remote end of the tunnel. It uses this information to look up the IP address of the remote end of the tunnel through the DDNS server associated with that domain name.

Local ID (peer ID)

The Local ID or peer ID can be used to uniquely identify one end of a VPN tunnel. This enables a more secure connection. Also if you have multiple VPN tunnels negotiating, this ensures the proper remote and local ends connect. When you configure it on your end, it is your Local ID. When the remote end connects to you, they see it as your peer ID.

If you are debugging a VPN connection, the Local ID is part of the VPN negotiations. You can use it to help troubleshoot connection problems.



In circumstances where multiple remote dialup VPN tunnels exist, each tunnel must have a peer ID set.

Configuring your Local ID

1. Go to **VPN > IPsec Wizard** and create the new custom tunnel or go to **VPN > IPsec Tunnels** and edit an existing tunnel.
2. Edit the **Phase 1 Proposal** (if it is not available, you may need to click the **Convert To Custom Tunnel** button).
3. In the **Phase 1 Proposal** section, enter your **Local ID**.
4. Select **OK**.

The default configuration is to accept all local IDs (peer IDs). If you have **Local ID** set, the remote end of the tunnel must be configured to accept your local ID.

Accepting a specific Peer ID

1. Go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
2. Edit **Authentication** (if it is not available, you may need to click the **Convert To Custom Tunnel** button).

3. Set **Mode** to **Aggressive**.
4. For **Peer Options**, select **This peer ID**. This option becomes visible only when **Aggressive** mode is selected.
5. In the **Peer ID** field, enter the string the other end of the tunnel used for its local ID.
6. Configure the rest of the Phase 1 entry as required.
7. Select **OK**.

Route-based or policy-based VPN

VPN over dynamic DNS can be configured with either route-based or policy-based VPN settings. Both are valid, but have differences in configuration. Choose the best method based on your requirements. For more information on route-based and policy-based, see [IPsec VPN overview on page 1638](#).

Route-based VPN configuration requires two security policies to be configured (one for each direction of traffic) to permit traffic over the VPN virtual interface, and you must also add a static route entry for that VPN interface or the VPN traffic will not reach its destination. See [Dynamic DNS configuration on page 1723](#) and [Dynamic DNS configuration on page 1723](#).

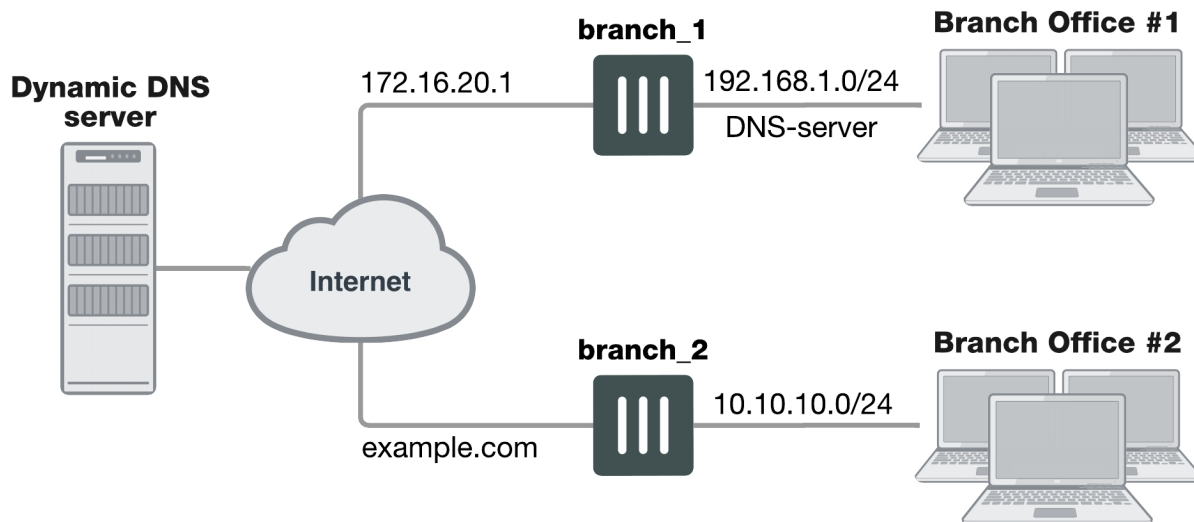
Policy-based VPN configuration uses more complex and often more IPsec security policies, but does not require a static route entry. It has the benefit of being able to configure multiple policies for handling multiple protocols in different ways, such as more scanning of less secure protocols or guaranteeing a minimum bandwidth for protocols such as VoIP. See [Dynamic DNS configuration on page 1723](#) and [Dynamic DNS configuration on page 1723](#).

DDNS topology

In this scenario, two branch offices each have a FortiGate unit and are connected in a gateway-to-gateway VPN configuration. One FortiGate unit has a domain name (example.com) with a dynamic IP address. See `branch_2` in the figure below.

Whenever the `branch_2` unit connects to the Internet (and possibly also at predefined intervals set by the ISP), the ISP may assign a different IP address to the FortiGate unit. The unit has its domain name registered with a dynamic DNS service. The `branch_2` unit checks in with the DDNS server on a regular basis, and that server provides the DNS information for the domain name, updating the IP address from time to time. Remote peers have to locate the `branch_2` FortiGate unit through a DNS lookup each time to ensure the address they get is current and correct.

Example dynamic DNS configuration



When a remote peer (such as the `branch_1` FortiGate unit above) initiates a connection to `example.com`, the local DNS server looks up and returns the IP address that matches the domain name `example.com`. The remote peer uses the retrieved IP address to establish a VPN connection with the `branch_2` FortiGate unit.

Assumptions

- You have administrator access to both FortiGate units.
- Both FortiGate units have interfaces named `wan1` and `internal`. (If not, you can use the alias feature to assign these labels as “nicknames” to other interfaces to follow this example.)
- Both FortiGate units have the most recent firmware installed, have been configured for their networks, and are currently passing normal network traffic.
- The `branch_2` FortiGate unit has its `wan1` interface defined as a dynamic DNS interface with the domain name of **example.com**.
- A basic gateway-to-gateway configuration is in place (see [Gateway-to-gateway configurations on page 1](#)) except one of the FortiGate units has a static domain name and a dynamic IP address instead of a static IP address.
- The FortiGate unit with the domain name is subscribed to one of the supported dynamic DNS services. Contact one of the services to set up an account. For more information and instructions about how to configure the FortiGate unit to push its dynamic IP address to a dynamic DNS server, see the [System Administration](#) handbook chapter.

Configuration overview

When a FortiGate unit receives a connection request from a remote VPN peer, it uses IPsec Phase 1 parameters to establish a secure connection and authenticate the VPN peer. Then, if the security policy permits the connection, the FortiGate unit establishes the tunnel using IPsec Phase 2 parameters and applies the security policy. Key management, authentication, and security services are negotiated dynamically through the IKE protocol.

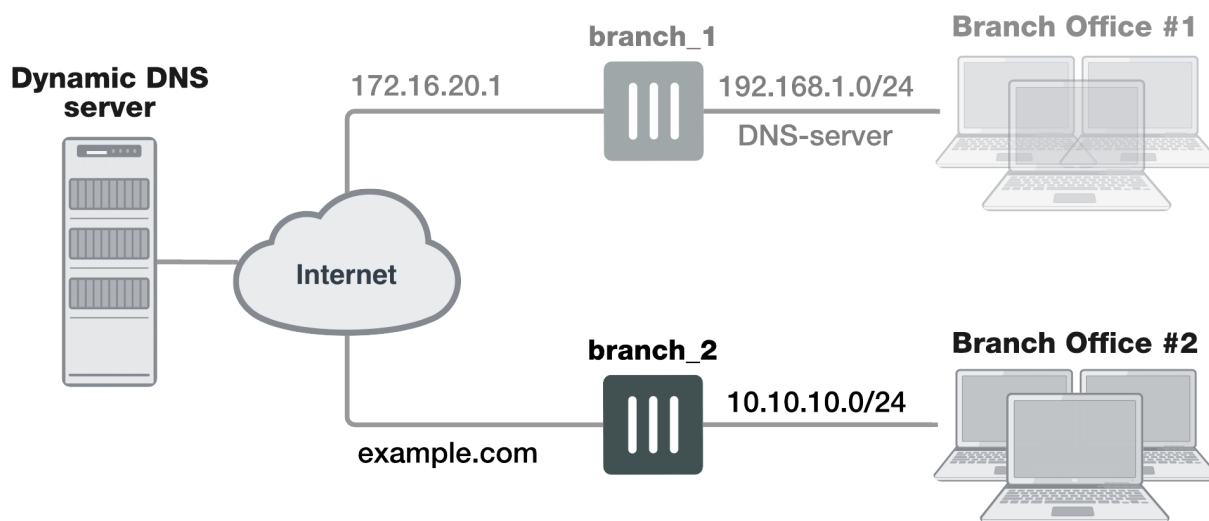
To support these functions, the following general configuration steps must be performed:

- Configure the branch_2 FortiGate unit with the dynamic IP address. This unit uses a Local ID string instead of an IP address to identify itself to the remote peer. See [Configuring the dynamically-addressed VPN peer](#) below, which is made up of configuring branch_2's VPN tunnel settings and security policies.
- Configure the fixed-address VPN peer. To initiate a VPN tunnel with the dynamically-addressed peer, this unit must first retrieve the IP address for the domain from the dynamic DNS service. See [Configuring the fixed-address VPN peer](#), which is made up of configuring branch_1's VPN tunnel settings and security policies.

Configuring the dynamically-addressed VPN peer

It is assumed that this FortiGate unit (branch_2) has already had its public facing interface, for example the wan1, configured with the proper dynamic DNS configuration.

Configuring branch_2, the dynamic address side



Define the Phase 1 parameters needed to establish a secure connection with the remote peer. See [Phase 1 parameters on page 1655](#). During this procedure you need to choose if you will be using route-based or policy-based VPNs.

1. Go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
2. Edit **Network** (full configuration options are only available once you click the **Convert To Custom Tunnel** button).
3. Enter the following information:

Remote Gateway	<p>Select Static IP Address.</p> <p>The remote peer this FortiGate is connecting to has a static IP public address.</p> <p>If the remote interface is PPPoE do not select Retrieve default gateway from server.</p>
IP Address	Enter 172.16.20.1, the IP address of the public interface to the remote peer.

Interface	Select the Internet-facing interface wan1 (selected by default).
NAT Traversal	Select Enable (selected by default).
Keepalive Frequency	Enter a keepalive frequency (In seconds; set to 10 by default).
Dead Peer Detection	Select a dead peer detection option. On Idle will attempt to reestablish VPN tunnels when a connection becomes idle (the idle interval is not a negotiated value). Use of periodic dead peer detection incurs extra overhead. When communicating to large numbers of IKE peers, you should consider using On Demand . (set to On Demand by default).

4. Edit **Authentication** and complete the following:

Mode	Select Aggressive .
-------------	----------------------------

5. Edit **Phase 1 Proposal** and complete the following:

Local ID	Enter <code>example.com</code> . A character string used by the <code>branch_2</code> FortiGate unit to identify itself to the remote peer. This value must be identical to the value in the This peer ID field of the Phase 1 remote gateway configuration on the <code>branch_1</code> remote peer. See Configuration overview on page 1726 .
-----------------	--

6. Open the **Phase 2 Selectors** panel.
Define the Phase 2 parameters needed to create a VPN tunnel with the remote peer. For details on Phase 2, see [Phase 2 parameters on page 1675](#).
7. Enter the following information and select **OK**.

Name	Automatically entered as the name of the VPN tunnel.
Phase 1	Select <code>branch_2</code> . The name of the Phase 1 configuration that you defined earlier.

Define security policies to permit communications between the private networks through the VPN tunnel. Route-based and policy-based VPNs require different security policies. For detailed information about creating security policies, see [Defining VPN security policies on page 1](#).

After defining the two address ranges, select one of [Creating branch_2 route-based security policies on page 1729](#) or [Creating branch_2 policy-based security policies on page 1731](#) to configure the appropriate VPN policies.

Define VPN connection names for the address ranges of the private networks. These addresses are used in the security policies that permit communication between the networks. For more information, see [Defining VPN security policies on page 1](#).

Define an address name for the IP address and netmask of the private network behind the local FortiGate unit.

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**.
3. Enter the following information, and select **OK**.

Name	Enter <code>branch_2_internal</code> . Enter a meaningful name.
Type	Select IP/Netmask .
Subnet / IP Range	Enter <code>10.10.10.0/24</code> . Include the netmask or specify a specific range.
Interface	Select internal . The interface that will be handling the traffic from the internal network.

Define an address name for the IP address and netmask of the private network behind the remote peer.

4. Select **Create New**.
5. Enter the following information, and select **OK**.

Name	Enter <code>branch_1_internal</code> . A meaningful name for the private network at the remote end of the VPN tunnel.
Type	Select IP/Netmask .
Subnet / IP Range	Enter <code>192.168.1.0/24</code> . Include the netmask. Optionally you can specify a range
Interface	Select any . The interface that will be handling the remote VPN traffic on this FortiGate unit. If you are unsure, or multiple interfaces may be handling this traffic use <code>any</code> .

Creating branch_2 route-based security policies

Define ACCEPT security policies to permit communication between the branch_2 and branch_1 private networks. Once the route-based policy is configured a routing entry must be configured to route traffic over the VPN interface.

Define a policy to permit the branch_2 local FortiGate unit to initiate a VPN session with the branch_1 VPN peer.

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information, and select **OK**.

Name	Enter an appropriate name for the policy.
-------------	---

Incoming Interface	Select internal . The interface that connects to the private network behind this FortiGate unit.
Outgoing Interface	Select branch_2 . The VPN Tunnel (IPsec Interface).
Source	Select branch_2_internal . Select the address name for the private network behind this FortiGate unit.
Destination Address	Select branch_1_internal . The address name the private network behind the remote peer.
Action	Select ACCEPT .
NAT	Disable NAT .
Comments	Route-based: Initiate a branch_2 to branch_1 VPN tunnel.

Define a policy to permit the branch_1 remote VPN peer to initiate VPN sessions.

3. Select **Create New**.
4. Enter the following information, and select **OK**.

Name	Enter an appropriate name for the policy.
Incoming Interface	Select branch_2 . The VPN Tunnel (IPsec Interface).
Outgoing Interface	Select internal . The interface connecting the private network behind this FortiGate unit.
Source	Select branch_1_internal . The address name for the private network behind the remote peer.
Destination Address	Select branch_2_internal . The address name for the private network behind this FortiGate unit.
Action	Select ACCEPT .
NAT	Disable NAT .
Comments	Route-based: Initiate a branch_1 to branch_2 internal VPN tunnel.

5. Optionally configure any other security policy settings you require such as UTM or traffic shaping for this policy.
6. Place these policies in the policy list above any other policies having similar source and destination addresses. This will ensure VPN traffic is matched against the VPN policies before any other policies.

Creating routing entry for VPN interface - CLI

```
config router static
edit 5
set dst 0.0.0.0 0.0.0.0
```

```

        set dynamic-gateway enable
        set device wan1
    next
end

```

This routing entry must be added in the CLI because the dynamic-gateway option is not available in the web-based manager.

Creating branch_2 policy-based security policies

Define an IPsec policy to permit VPN sessions between the private networks. Define an IPsec policy to permit the VPN sessions between the local branch_2 unit and the remote branch_1 unit.

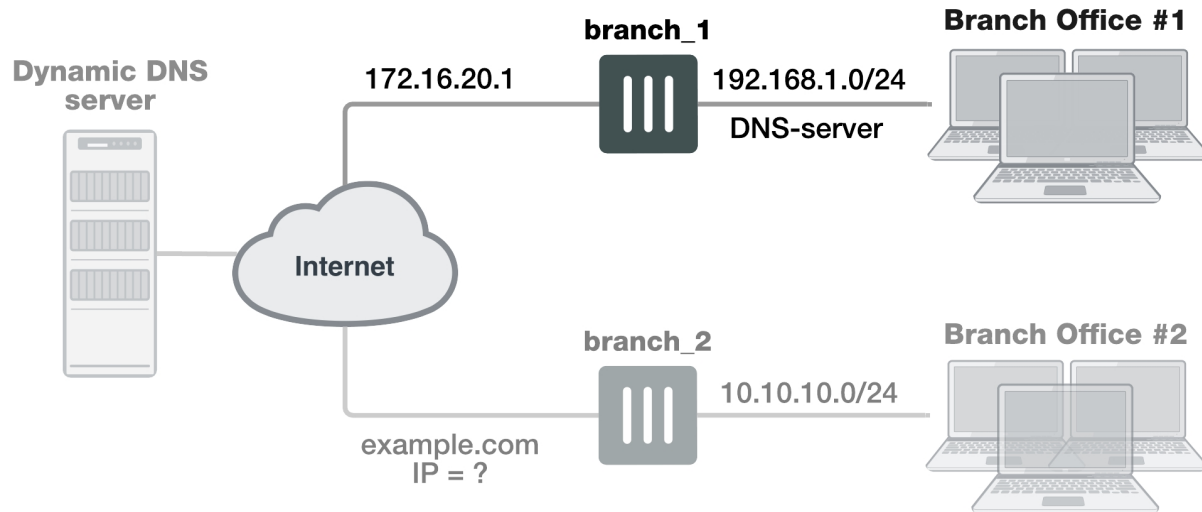
1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information, and select **OK**.

Name	Enter an appropriate name for the policy.
Incoming Interface	Select internal . The interface connecting the private network behind this FortiGate unit.
Outgoing Interface	Select wan1 . The FortiGate unit's public interface.
Source	Select branch_2_internal . The address name for the private network behind this local FortiGate unit.
Destination Address	Select branch_1_internal . The address name for the private network behind branch_1, the remote peer.
Action	Select IPsec . Under VPN Tunnel , select branch_2 from the drop-down list. The name of the Phase 1 tunnel. Select Allow traffic to be initiated from the remote site .
Comments	Policy-based: allows traffic in either direction to initiate the VPN tunnel.

3. Optionally configure any other security policy settings you require such as UTM or traffic shaping for this policy.
4. Place these policies in the policy list above any other policies having similar source and destination addresses. This will ensure VPN traffic is matched against the VPN policies before any other policies.

Configuring the fixed-address VPN peer

The fixed-address VPN peer, branch_1, needs to retrieve the IP address from the dynamic DNS service to initiate communication with the dynamically-addressed peer, branch_2. It also depends on the peer ID (local ID) to initiate the VPN tunnel with branch_2.



Define the Phase 1 parameters needed to establish a secure connection with the remote peer. For more information, see [Phase 1 parameters on page 1655](#).

1. Go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
2. Edit **Network** (if it is not available, you may need to click the **Convert to Custom Tunnel** button).
3. Enter the following information and select **OK**.

Remote Gateway	Select Dynamic DNS . The remote peer this FortiGate is connecting to has a dynamic IP address.
Dynamic DNS	Type the fully qualified domain name of the remote peer (for example, <code>example.com</code>).
Interface	Select wan1 . The public facing interface on the fixed-address FortiGate unit.
Mode Config	Select Aggressive .
Peer Options	Select This peer ID , and enter <code>example.com</code> . This option only appears when the mode is set to Aggressive. The identifier of the FortiGate unit with the dynamic address.

4. Edit **Authentication**, enter the following information and select **OK**.

Peer Options	Select This peer ID , and enter <code>example.com</code> . This option only appears when the authentication method is set to Signature . The identifier of the FortiGate unit with the dynamic address.
---------------------	---

5. Define the Phase 2 parameters needed to create a VPN tunnel with the remote peer. See [Phase 2 parameters on page 1675](#). Enter these settings in particular:

Name	Enter <code>branch_1_p2</code> . A name to identify this Phase 2 configuration.
-------------	---

Phase 1Select **branch_1**.

The name of the Phase 1 configuration that you defined for the remote peer. You can select the name of the remote gateway from the Dynamic DNS part of the list.

The `branch_1` FortiGate unit has a fixed IP address and will be connecting to the `branch_2` FortiGate unit that has a dynamic IP address and a domain name of `example.com`. Remember if you are using route-based security policies that you must add a route for the VPN traffic.

Defining address ranges for branch_1 security policies

As with `branch_2` previously, `branch_1` needs address ranges defined as well. See [Defining policy addresses on page 1](#).

1. Go to **Policy & Objects > Addresses** and select **Create New > Address**.
2. Enter the following information, and select **OK**.

Name	Enter <code>branch_2_internal</code> . A meaningful name for the private network behind the <code>branch_2</code> FortiGate unit.
Type	Select IP/Netmask .
Subnet / IP Range	Enter <code>10.10.10.0/24</code> . Include the netmask or specify a specific range.
Interface	Select internal . This is the interface on this FortiGate unit that will be handling with this traffic.

3. Define an address name for the IP address and netmask of the private network behind the remote peer.
4. Create another address. Enter the following information, and select **OK**.

Name	Enter <code>branch_1_internal</code> . A meaningful name for the private network behind the <code>branch_1</code> peer.
Type	Select IP/Netmask .
Subnet / IP Range	Enter <code>192.168.1.0/24</code> . Include the netmask or specify a specific range.
Interface	Select any . The interface on this FortiGate unit that will be handling with this traffic. If you are unsure, or multiple interfaces may be handling this traffic use <code>any</code> .

Creating branch_1 route-based security policies

Define an ACCEPT security policy to permit communications between the source and destination addresses. See [Defining VPN security policies on page 1](#).

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information, and select **OK**.

Name	Enter an appropriate name for the policy.
Incoming Interface	Select internal . The interface that connects to the private network behind the <code>branch_1</code> FortiGate unit.
Outgoing Interface	Select branch_1 . The VPN Tunnel (IPsec Interface) you configured earlier.
Source	Select branch_1_internal . The address name that you defined for the private network behind this FortiGate unit.
Destination Address	Select branch_2_internal . The address name that you defined for the private network behind the <code>branch_2</code> peer.
Action	Select ACCEPT .
NAT	Disable NAT .
Comments	Internal -> branch2

To permit the remote client to initiate communication, you need to define a security policy for communication in that direction.

3. Select **Create New**.
4. Enter the following information, and select **OK**.

Name	Enter an appropriate name for the policy.
Incoming Interface	Select branch_1 . The VPN Tunnel (IPsec Interface) you configured earlier.
Outgoing Interface	Select internal . The interface that connects to the private network behind this FortiGate unit.
Source	Select branch_2_internal . The address name that you defined for the private network behind the <code>branch_2</code> remote peer.
Destination Address	Select branch_1_internal . The address name that you defined for the private network behind this FortiGate unit.
Action	Select ACCEPT .
NAT	Disable NAT .
Comments	branch_2 -> Internal

Creating branch_1 policy-based security policies

A policy-based security policy allows you the flexibility to allow inbound or outbound traffic or both through this single policy.

This policy-based IPsec VPN security policy allows both inbound and outbound traffic

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information, and select **OK**.

Incoming Interface	Select internal . The interface that connects to the private network behind this FortiGate unit.
Outgoing Interface	Select wan1 . The FortiGate unit's public interface.
Source	Select branch_1_internal . The address name that you defined for the private network behind this FortiGate unit.
Destination Address	Select branch_2_internal . The address name that you defined for the private network behind the remote peer.
Action	Select IPsec . Under VPN Tunnel , select branch_1 from the drop-down list. The name of the Phase 1 tunnel. Select Allow traffic to be initiated from the remote site .

- Place this security policy in the policy list above any other policies having similar source and destination addresses.

Results

Once both ends are configured, you can test the VPN tunnel.

To test the VPN initiated by branch_2

- On branch_2, go to **Monitor > IPsec Monitor**.
All IPsec VPN tunnels will be listed on this page, no matter if they are connected or disconnected.
- Select the tunnel listed for branch_2, and select the status column for that entry.
The status will say **Bring Up** and remote port, incoming and outgoing data will all be zero. This indicates an inactive tunnel. When you right-click and select **Bring Up**, the FortiGate will try to set up a VPN session over this tunnel. If it is successful, Bring Up will change to Active, and the arrow icon will change to a green up arrow icon.
- If this does not create a VPN tunnel with increasing values for incoming and outgoing data, you need to start troubleshooting:

To test the VPN initiated by branch_1

- On branch_1, go to **Monitor > IPsec Monitor**.
- Select the tunnel listed for branch_1, and select the status column.
The difference between branch_2 and branch_1 at this point is that the tunnel entry for branch-1 will not have a remote gateway IP address. It will be resolved when the VPN tunnel is started.
- If this does not create a VPN tunnel with increasing values for incoming and outgoing data, you need to start troubleshooting.

Some troubleshooting ideas include:

- If there was no entry for the tunnel on the monitor page, check the Auto Key (IKE) page to verify the Phase 1 and Phase 2 entries exist.
- Check the security policy or policies, and ensure there is an outgoing policy as a minimum.
- Check that you entered a local ID in the Phase 1 configuration, and that branch_1 has the same local ID.
- Ensure the local DNS server has an up-to-date DNS entry for exmaple.com.

For more information, see [Troubleshooting on page 1](#).

FortiClient dialup-client configuration

The FortiClient Endpoint Security application is an IPsec VPN client with antivirus, antispam and firewall capabilities. This section explains how to configure dialup VPN connections between a FortiGate unit and one or more FortiClient Endpoint Security applications.

FortiClient users are usually mobile or remote users who need to connect to a private network behind a FortiGate unit. For example, the users might be employees who connect to the office network while traveling or from their homes.

For greatest ease of use, the FortiClient application can download the VPN settings from the FortiGate unit to configure itself automatically.

The following topics are included in this section:

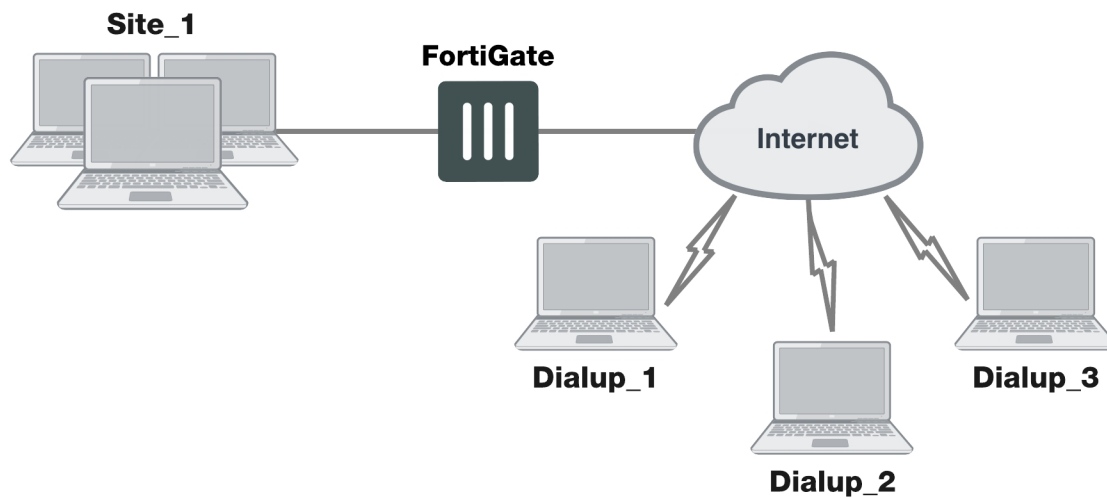
Configuration overview

Dialup users typically obtain dynamic IP addresses from an ISP through Dynamic Host Configuration Protocol (DHCP) or Point-to-Point Protocol over Ethernet (PPPoE). Then, the FortiClient Endpoint Security application initiates a connection to a FortiGate dialup server.

By default the FortiClient dialup client has the same IP address as the host PC on which it runs. If the host connects directly to the Internet, this is a public IP address. If the host is behind a NAT device, such as a router, the IP address is a private IP address. The NAT device must be NAT traversal (NAT-T) compatible to pass encrypted packets (see [Phase 1 parameters on page 1655](#)). The FortiClient application also can be configured to use a virtual IP address (VIP). For the duration of the connection, the FortiClient application and the FortiGate unit both use the VIP address as the IP address of the FortiClient dialup client.

The FortiClient application sends its encrypted packets to the VPN remote gateway, which is usually the public interface of the FortiGate unit. It also uses this interface to download VPN settings from the FortiGate unit. See [Automatic configuration of FortiClient dialup clients on page 1737](#).

Example FortiClient dialup-client configuration



Peer identification

The FortiClient application can establish an IPsec tunnel with a FortiGate unit configured to act as a dialup server. When the FortiGate unit acts as a dialup server, it does not identify the client using the Phase 1 remote gateway address. The IPsec tunnel is established if authentication is successful and the IPsec security policy associated with the tunnel permits access. If configured, the FortiGate unit could also require FortiClient registration, that is, the remote user would be required to have FortiClient installed before connection is completed.

Automatic configuration of FortiClient dialup clients

The FortiClient application can obtain its VPN settings from the FortiGate VPN server. FortiClient users need to know only the FortiGate VPN server IP address and their username and password on the FortiGate unit.

The FortiGate unit listens for VPN policy requests from clients on TCP port 8900. When the dialup client connects:

- The client initiates a Secure Sockets Layer (SSL) connection to the FortiGate unit.
- The FortiGate unit requests a user name and password from the FortiClient user. Using these credentials, it authenticates the client and determines which VPN policy applies to the client.
- Provided that authentication is successful, the FortiGate unit downloads a VPN policy to the client over the SSL connection. The information includes IPsec Phase 1 and Phase 2 settings, and the IP addresses of the private networks that the client is authorized to access.
- The client uses the VPN policy settings to establish an IPsec Phase 1 connection and Phase 2 tunnel with the FortiGate unit.

FortiClient-to-FortiGate VPN configuration steps

Configuring dialup client capability for FortiClient dialup clients involves the following general configuration steps:

1. If you will be using VIP addresses to identify dialup clients, determine which VIP addresses to use. As a precaution, consider using VIP addresses that are not commonly used.
2. Configure the FortiGate unit to act as a dialup server. See [Configure the FortiGate unit on page 1](#).

3. If the dialup clients will be configured to obtain VIP addresses through DHCP over IPsec, configure the FortiGate unit to act as a DHCP server or to relay DHCP requests to an external DHCP server.
4. Configure the dialup clients. See [Configure the FortiClient Endpoint Security application on page 1](#).

Using virtual IP addresses

When the FortiClient host PC is located behind a NAT device, unintended IP address overlap issues may arise between the private networks at the two ends of the tunnel. For example, the client's host might receive a private IP address from a DHCP server on its network that by co-incidence is the same as a private IP address on the network behind the FortiGate unit. A conflict will occur in the host's routing table and the FortiClient Endpoint Security application will be unable to send traffic through the tunnel. Configuring virtual IP (VIP) addresses for FortiClient applications prevents this problem.

Using VIPs ensures that client IP addresses are in a predictable range. You can then define security policies that allow access only to that source address range. If you do not use VIPs, the security policies must allow all source addresses because you cannot predict the IP address for a remote mobile user.

The FortiClient application must not have the same IP address as any host on the private network behind the FortiGate unit or any other connected FortiClient application. You can ensure this by reserving a range of IP addresses on the private network for FortiClient users. Or, you can assign FortiClient VIPs from an uncommonly used subnet such as 10.254.254.0/24 or 192.168.254.0/24.

You can reserve a VIP address for a particular client according to its device MAC address and type of connection. The DHCP server then always assigns the reserved VIP address to the client. For more information about this feature, see the "dhcp reserved-address" section in the "system" chapter of the [FortiGate CLI Reference](#).



On the host computer, you can find out the VIP address that the FortiClient Endpoint Security application is using. For example, in Windows command prompt, type `ipconfig /all`

On Linux or Mac OS X, type `ifconfig` in a terminal window. The output will also show the IP address that has been assigned to the host Network Interface Card (NIC).

It is best to assign VIPs using DHCP over IPsec. The FortiGate dialup server can act as a DHCP server or relay requests to an external DHCP server. You can also configure VIPs manually on FortiClient applications, but it is more difficult to ensure that all clients use unique addresses.

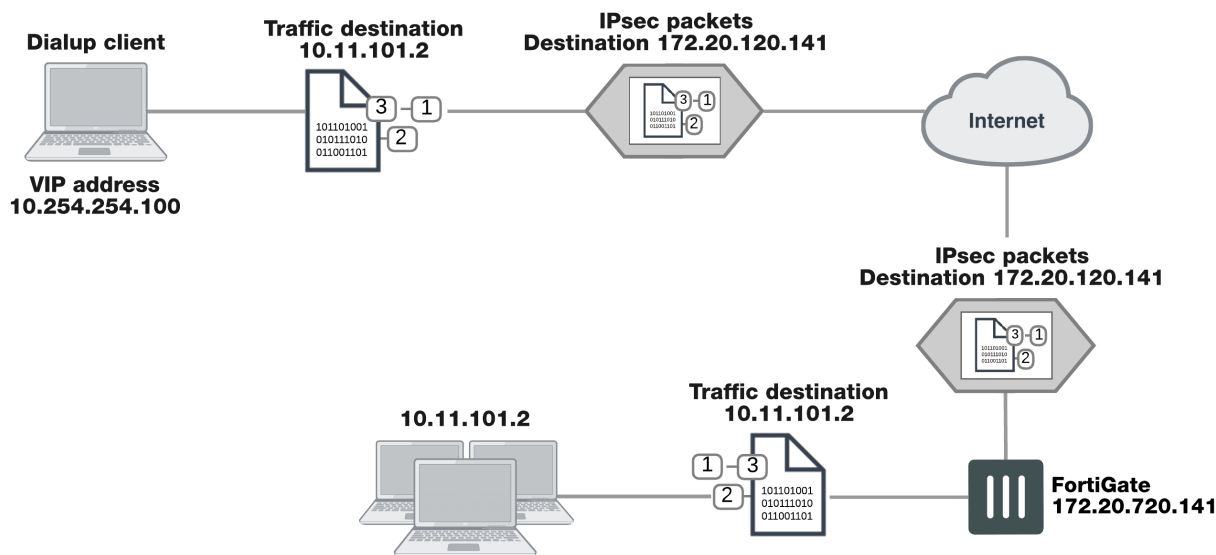


If you assign a VIP on the private network behind the FortiGate unit and enable DHCP-IPsec (a Phase 2 advanced option), the FortiGate unit acts as a proxy on the local private network for the FortiClient dialup client. Whenever a host on the network behind the dialup server issues an ARP request for the device MAC address of the FortiClient host, the FortiGate unit answers the ARP request on behalf of the FortiClient host and forwards the associated traffic to the FortiClient host through the tunnel. For more information, see [Phase 2 parameters on page 1675](#).

FortiGate units fully support [RFC 3456](#). The FortiGate DHCP over IPsec feature can be enabled to allocate VIP addresses to FortiClient dialup clients using a FortiGate DHCP server.

The figure below shows an example of a FortiClient-to-FortiGate VPN where the FortiClient application is assigned a VIP on an uncommonly used subnet. The diagram also shows that while the destination for the information in the encrypted packets is the private network behind the FortiGate unit, the destination of the IPsec packets themselves is the public interface of the FortiGate unit that acts as the end of the VPN tunnel.

IP address assignments in a FortiClient dialup-client configuration



Assigning VIPs by RADIUS user group

If you use XAuth authentication, you can assign users the virtual IP address stored in the Framed-IP-Address field of their record on the RADIUS server. (See [RFC 2865](#) and [RFC 2866](#) for more information about RADIUS fields.) To do this:

- Set the DHCP server **IP Assignment Mode** to **User-group defined method**. This is an Advanced setting. See [Configuring a DHCP server on a FortiGate interface on page 1743](#).
- Create a new firewall user group and add the RADIUS server to it.
- In your Phase 1 settings, configure the FortiGate unit as an XAuth server and select from **User Group** the new user group that you created. For more information, see [Phase 1 parameters on page 1655](#).
- Configure the FortiClient application to use XAuth. See [Configuration overview on page 1736](#).

FortiClient dialup-client infrastructure requirements

- To support policy-based VPNs, the FortiGate dialup server may operate in either NAT mode or transparent mode. NAT mode is required if you want to create a route-based VPN.
- If the FortiClient dialup clients will be configured to obtain VIP addresses through FortiGate DHCP relay, a DHCP server must be available on the network behind the FortiGate unit and the DHCP server must have a direct route to the FortiGate unit.
- If the FortiGate interface to the private network is not the default gateway, the private network behind the FortiGate unit must be configured to route IP traffic destined for dialup clients back (through an appropriate gateway) to the FortiGate interface to the private network. As an alternative, you can configure the IPsec security policy on the FortiGate unit to perform inbound NAT on IP packets. Inbound NAT translates the source addresses of inbound decrypted packets into the IP address of the FortiGate interface to the local private network.

Configuring the FortiGate unit

Configuring the FortiGate unit to establish VPN connections with FortiClient Endpoint Security users involves the following steps:

- Configure the VPN settings
- If the dialup clients use automatic configuration, configure the FortiGate unit as a VPN policy server
- If the dialup clients obtain VIP addresses by DHCP over IPsec, configure an IPsec DHCP server or relay

The procedures in this section cover basic setup of policy-based and route-based VPNs compatible with FortiClient Endpoint Security. A route-based VPN is simpler to configure.



The IPsec VPN Wizard greatly simplifies IPsec VPN tunnel creation for route-based tunnels.

To configure FortiGate unit VPN settings to support FortiClient users, you need to:

- Configure the FortiGate Phase 1 VPN settings
- Configure the FortiGate Phase 2 VPN settings
- Add the security policy

On the local FortiGate unit, define the Phase 1 configuration needed to establish a secure connection with the FortiClient peer. See [Phase 1 parameters on page 1655](#).

1. Go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
2. Edit **Network** (full configuration options are only available once you click the **Convert To Custom Tunnel** button).
3. Enter these settings in particular:

Remote Gateway	Select Dialup User .
IP Address	Enter the IP address of the remote peer.
Interface	Select the interface through which clients connect to the FortiGate unit.
Mode Config	When enabled, further options become available: <ul style="list-style-type: none"> • Client Address Range • Subnet Mask • Use System DNS • DNS Server • Enable IPv4 Split Tunnel
Authentication Method	Select Pre-shared Key .
Pre-shared Key	Enter the pre-shared key. This must be the same preshared key provided to the FortiClient users.
Peer option	Select Any peer ID .

4. Edit **Authentication** and enter the following information:

Method	Select Pre-shared Key .
Pre-shared Key	Enter the pre-shared key. This must be the same preshared key provided to the FortiClient users.
Peer Options	Set Accept Types to Any peer ID .

5. Define the Phase 2 parameters needed to create a VPN tunnel with the FortiClient peer. See [Phase 2 parameters on page 1675](#). Enter these settings in particular:

Name	Enter a name to identify this Phase 2 configuration.
Phase 1	Select the name of the Phase 1 configuration that you defined.
Advanced	Select to configure the following optional setting.
DHCP-IPsec	Select if you provide virtual IP addresses to clients using DHCP.

6. Define names for the addresses or address ranges of the private networks that the VPN links. These addresses are used in the security policies that permit communication between the networks. For more information, see [Defining policy addresses on page 1](#).

Enter these settings in particular:

- Define an address name for the individual address or the subnet address that the dialup users access through the VPN.
 - If FortiClient users are assigned VIP addresses, define an address name for the subnet to which these VIPs belong.
4. Define security policies to permit communication between the private networks through the VPN tunnel. Route-based and policy-based VPNs require different security policies. For detailed information about creating security policies, see [Defining VPN security policies on page 1](#).

If the security policy, which grants the VPN Connection is limited to certain services, DHCP must be included, otherwise the client won't be able to retrieve a lease from the FortiGate's (IPsec) DHCP server, because the DHCP Request (coming out of the tunnel) will be blocked.

Route-based VPN security policies

Define an ACCEPT security policy to permit communications between the source and destination addresses.

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter these settings in particular:

Name	Enter an appropriate name for the policy.
Incoming Interface	Select the VPN Tunnel (IPsec Interface) you configured in Step "Configuration overview" on page 1736 .
Outgoing Interface	Select the interface that connects to the private network behind this FortiGate unit.

Source	Select all .
Destination Address	Select all .
Action	Select ACCEPT .
NAT	Disable NAT .

If you want to allow hosts on the private network to initiate communications with the FortiClient users after the tunnel is established, you need to define a security policy for communication in that direction.

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter these settings in particular:

Incoming Interface	Select the interface that connects to the private network behind this FortiGate unit.
Outgoing Interface	Select the interface that connects to the private network behind this FortiGate unit.
Source	Select all .
Destination Address	Select all .
Action	Select ACCEPT .
NAT	Disable NAT .

Policy-based VPN security policy

Define an IPsec security policy to permit communications between the source and destination addresses.

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter these settings in particular:

Incoming Interface	Select the interface that connects to the private network behind this FortiGate unit.
Outgoing Interface	Select the FortiGate unit's public interface.
Source	Select the address name that you defined in Step "Configuration overview" on page 1736 for the private network behind this FortiGate unit.
Destination Address	If FortiClient users are assigned VIPs, select the address name that you defined for the VIP subnet. Otherwise, select all .
Action	Select IPsec . Under VPN Tunnel , select the name of the Phase 1 configuration that you created in Step "Configuration overview" on page 1736 from the drop-down list. Select Allow traffic to be initiated from the remote site .

Place VPN policies in the policy list above any other policies having similar source and destination addresses.

Configuring the FortiGate unit as a VPN policy server

When a FortiClient application set to automatic configuration connects to the FortiGate unit, the FortiGate unit requests a user name and password. If the user supplies valid credentials, the FortiGate unit downloads the VPN settings to the FortiClient application.

You must do the following to configure the FortiGate unit to work as a VPN policy server for FortiClient automatic configuration:

1. Create user accounts for FortiClient users.
2. Create a user group for FortiClient users and the user accounts that you created in step 1.
3. Connect to the FortiGate unit CLI and configure VPN policy distribution as follows:

```
config vpn ipsec forticlient
  edit <policy_name>
    set phase2name <tunnel_name>
    set usergroupname <group_name>
    set status enable
  end
```

<tunnel_name> must be the Name you specified in the step 2 of [Configuration overview on page 1736](#).

<group_name> must be the name of the user group your created for FortiClient users.

Configuring DHCP services on a FortiGate interface

If the FortiClient dialup clients are configured to obtain a VIP address using DHCP, configure the FortiGate dialup server to either:

- Relay DHCP requests to a DHCP server behind the FortiGate unit (see [Configuring DHCP relay on a FortiGate interface on page 1743](#) below).
- Act as a DHCP server (see [Configuring a DHCP server on a FortiGate interface on page 1743](#)).

Note that DHCP services are typically configured during the interface creation stage, but you can return to an interface to modify DHCP settings if need be.

Configuring DHCP relay on a FortiGate interface

1. Go to **Network > Interfaces** and select the interface that you want to relay DHCP.
2. Enable **DHCP Server**, and create a new DHCP **Address Range** and **Netmask**.
3. Open the **Advanced...** menu and set **Mode** to **Relay**.
4. Enter the **DHCP Server IP**.
5. Select **OK**.

Configuring a DHCP server on a FortiGate interface

1. Go to **Network > Interfaces** and select the interface that you want to act as a DHCP server.
2. Enable **DHCP Server**, and create a new DHCP **Address Range** and **Netmask**.
3. Set **Default Gateway** to **Specify**, and enter the IP address of the default gateway that the DHCP server assigns to DHCP clients.
4. Set **DNS Server** to **Same as System DNS**. If you want to use a different DNS server for VPN clients, select **Specify** and enter an IP address in the available field.

5. Open the **Advanced...** menu and set **Mode** to **Server**.
6. Select **OK**.

Configure the FortiClient Endpoint Security application

The following procedure explains how to configure the FortiClient Endpoint Security application to communicate with a remote FortiGate dialup server using the VIP address that you specify manually. These procedures are based on FortiClient 5.4.1.

Configuring FortiClient

This procedure explains how to configure the FortiClient application manually using the default IKE and IPsec settings. For more information, refer to the FortiClient Administration Guide.

1. Go to **Remote Access** and select the **Settings** icon.
2. Select **Add a new connection**, set the new VPN connection to **IPsec VPN**, and complete following information:

Connection Name	Enter a descriptive name for the connection.
Remote Gateway	Enter the IP address or the fully qualified domain name (FQDN) of the remote gateway.
Authentication Method	Select Pre-shared Key and enter the pre-shared key in the field provided.
Authentication (XAuth)	<p>Extended Authentication (XAuth) increases security by requiring additional user authentication in a separate exchange at the end of the VPN Phase 1 negotiation. The FortiGate unit challenges the user for a user name and password. It then forwards the user's credentials to an external RADIUS or LDAP server for verification.</p> <p>Implementation of XAuth requires configuration at both the FortiGate unit and the FortiClient application.</p>

3. Select **OK**.

Adding XAuth authentication

For information about configuring a FortiGate unit as an XAuth server, see [Phase 1 parameters on page 1655](#). The following procedure explains how to configure the FortiClient application.

Note that XAuth is not compatible with IKE version 2.

For more information on configuring XAuth authentication, see the [FortiClient Administration Guide](#).

FortiGate dialup-client configurations


This section explains how to set up a FortiGate dialup-client IPsec VPN. In a FortiGate dialup-client configuration, a FortiGate unit with a static IP address acts as a dialup server and a FortiGate unit having a dynamic IP address initiates a VPN tunnel with the FortiGate dialup server.

The following topics are included in this section:

Configuration overview

A dialup client can be a FortiGate unit. The FortiGate dialup client typically obtains a dynamic IP address from an ISP through the Dynamic Host Configuration Protocol (DHCP) or Point-to-Point Protocol over Ethernet (PPPoE) before initiating a connection to a FortiGate dialup server.

Example FortiGate dialup-client configuration



In a dialup-client configuration, the FortiGate dialup server does not rely on a Phase 1 remote gateway address to establish an IPsec VPN connection with dialup clients. As long as authentication is successful and the IPsec security policy associated with the tunnel permits access, the tunnel is established.

Several different ways to authenticate dialup clients and restrict access to private networks based on client credentials are available. To authenticate FortiGate dialup clients and help to distinguish them from FortiClient dialup clients when multiple clients will be connecting to the VPN through the same tunnel, best practices dictate that you assign a unique identifier (local ID or peer ID) to each FortiGate dialup client. For more information, see [Phase 1 parameters on page 1655](#).



Whenever you add a unique identifier (local ID) to a FortiGate dialup client for identification purposes, you must select Aggressive mode on the FortiGate dialup server and also specify the identifier as a peer ID on the FortiGate dialup server. For more information, see [Phase 1 parameters on page 1655](#).

Users behind the FortiGate dialup server cannot initiate the tunnel because the FortiGate dialup client does not have a static IP address. After the tunnel is initiated by users behind the FortiGate dialup client, traffic from the private network behind the FortiGate dialup server can be sent to the private network behind the FortiGate dialup client.

Encrypted packets from the FortiGate dialup client are addressed to the public interface of the dialup server. Encrypted packets from the dialup server are addressed either to the public IP address of the FortiGate dialup client (if the dialup client connects to the Internet directly), or if the FortiGate dialup client is behind a NAT device, encrypted packets from the dialup server are addressed to the public IP address of the NAT device.

If a router with NAT capabilities is in front of the FortiGate dialup client, the router must be NAT-T compatible for encrypted traffic to pass through the NAT device. For more information, see [Phase 1 parameters on page 1655](#).

When the FortiGate dialup server decrypts a packet from the FortiGate dialup client, the source address in the IP header may be one of the following values, depending on the configuration of the network at the far end of the tunnel:

- If the FortiGate dialup client connects to the Internet directly, the source address will be the private IP address of a host or server on the network behind the FortiGate dialup client.
- If the FortiGate dialup client is behind a NAT device, the source address will be the public IP address of the NAT device.

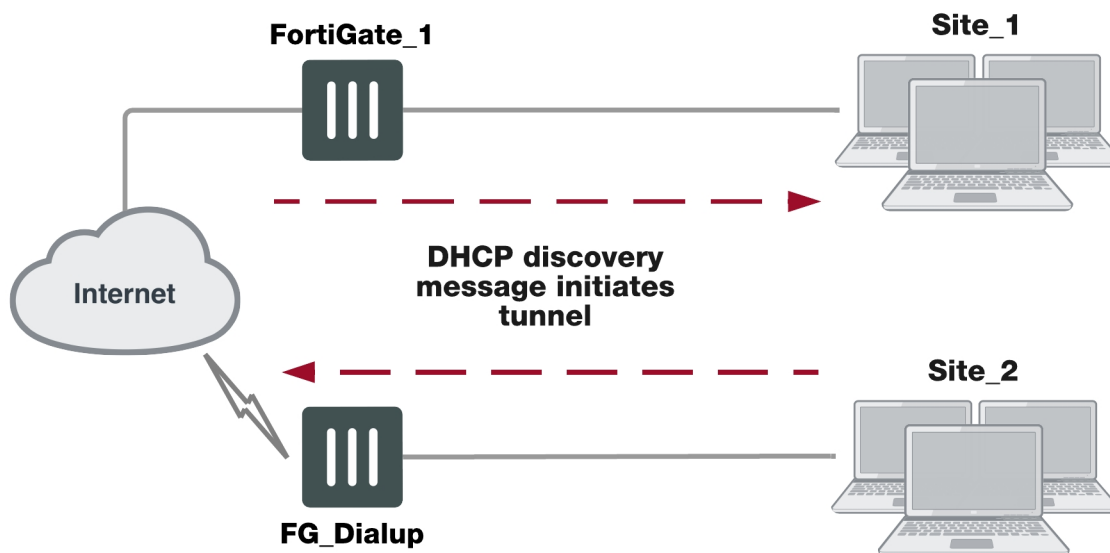
In some cases, computers on the private network behind the FortiGate dialup client may (by co-incidence) have IP addresses that are already used by computers on the network behind the FortiGate dialup server. In this type of situation (ambiguous routing), conflicts may occur in one or both of the FortiGate routing tables and traffic destined for the remote network through the tunnel may not be sent.

In many cases, computers on the private network behind the FortiGate dialup client will most likely obtain IP addresses from a local DHCP server behind the FortiGate dialup client. However, unless the local and remote networks use different private network address spaces, unintended ambiguous routing and IP-address overlap issues may arise.

To avoid these issues, you can configure FortiGate DHCP relay on the dialup client instead of using a DHCP server on the network behind the dialup client. The FortiGate dialup client can be configured to relay DHCP requests from the local private network to a DHCP server that resides on the network behind the FortiGate dialup server. You configure the FortiGate dialup client to pass traffic from the local private network to the remote network by enabling FortiGate DHCP relay on the FortiGate dialup client interface that is connected to the local private network.

Afterward, when a computer on the network behind the dialup client broadcasts a DHCP request, the dialup client relays the message through the tunnel to the remote DHCP server. The remote DHCP server responds with a private IP address for the computer. To avoid ambiguous routing and network overlap issues, the IP addresses assigned to computers behind the dialup client cannot match the network address space used by the private network behind the FortiGate dialup server.

Preventing network overlap in a FortiGate dialup-client configuration



When the DHCP server resides on the private network behind the FortiGate dialup server, the IP destination address specified in the IPsec security policy on the FortiGate dialup client must refer to that network.



You must add a static route to the DHCP server FortiGate unit if it is not directly connected to the private network behind the FortiGate dialup server; its IP address does not match the IP address of the private network. Also, the destination address in the IPsec security policy on the FortiGate dialup client must refer to the DHCP server address. The DHCP server must be configured to assign a range of IP addresses different from the DHCP server's local network, and also different from the private network addresses behind the FortiGate dialup server. See [Routing on page 1](#).

FortiGate dialup-client infrastructure requirements

The requirements are:

- The FortiGate dialup server must have a static public IP address.
- NAT mode is required if you want to create a route-based VPN.
- The FortiGate dialup server may operate in either NAT mode or transparent mode to support a policy-based VPN.
- Computers on the private network behind the FortiGate dialup client can obtain IP addresses either from a DHCP server behind the FortiGate dialup client, or a DHCP server behind the FortiGate dialup server.
 - If the DHCP server resides on the network behind the dialup client, the DHCP server must be configured to assign IP addresses that do not match the private network behind the FortiGate dialup server.
 - If the DHCP server resides on the network behind the FortiGate dialup server, the DHCP server must be configured to assign IP addresses that do not match the private network behind the FortiGate dialup client.

Configuring the server to accept FortiGate dialup-client connections

The procedures in this section assume that computers on the private network behind the FortiGate dialup client obtain IP addresses from a local DHCP server. The assigned IP addresses do not match the private network

behind the FortiGate dialup server.



In situations where IP-address overlap between the local and remote private networks is likely to occur, FortiGate DHCP relay can be configured on the FortiGate dialup client to relay DHCP requests to a DHCP server behind the FortiGate dialup server. For more information, see [To configure DHCP relay on a FortiGate interface on page 1](#).

Configuring dialup client capability for FortiGate dialup clients involves the following general configuration steps:

- Determine which IP addresses to assign to the private network behind the FortiGate dialup client, and add the IP addresses to the DHCP server behind the FortiGate dialup client. Refer to the software supplier's documentation to configure the DHCP server.
- Configure the FortiGate dialup server. See [Configuration overview on page 1745](#).
- Configure the FortiGate dialup client. See [Configuration overview on page 1745](#).

Before you begin, optionally reserve a unique identifier (peer ID) for the FortiGate dialup client. The dialup client will supply this value to the FortiGate dialup server for authentication purposes during the IPsec Phase 1 exchange. In addition, the value will enable you to distinguish FortiGate dialup-client connections from FortiClient dialup-client connections. The same value must be specified on the dialup server and on the dialup client.



In circumstances where multiple remote dialup VPN tunnels exist, each tunnel must have a peer ID set.

At the FortiGate dialup server, define the Phase 1 parameters needed to authenticate the FortiGate dialup client and establish a secure connection. See [Phase 1 parameters on page 1655](#).

1. Go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
2. Edit **Network** (full configuration options are only available once you click the **Convert To Custom Tunnel** button).
3. Enter these settings in particular:

Remote Gateway	Select Dialup User .
Interface	Select the interface through which clients connect to the FortiGate unit.

4. Edit **Authentication** and enter the following information:

Mode	If you will be assigning an ID to the FortiGate dialup client, select Aggressive .
Peer Options	If you will be assigning an ID to the FortiGate dialup client, set Accept Types to This peer ID and type the identifier that you reserved for the FortiGate dialup client into the adjacent field.

5. Define the Phase 2 parameters needed to create a VPN tunnel with the FortiGate dialup client. See [Phase 2 parameters on page 1675](#). Enter these settings in particular:

Name	Enter a name to identify this Phase 2 configuration.
Phase 1	Select the name of the Phase 1 configuration that you defined.

- Define names for the addresses or address ranges of the private networks that the VPN links. See [Defining policy addresses on page 1](#). Enter these settings in particular:
 - Define an address name for the server, host, or network behind the FortiGate dialup server.
 - Define an address name for the private network behind the FortiGate dialup client.
- Define the security policies to permit communications between the private networks through the VPN tunnel. Route-based and policy-based VPNs require different security policies. For detailed information about creating security policies, see [Defining VPN security policies on page 1](#).

Route-based VPN security policy

Define an ACCEPT security policy to permit communications between hosts on the private network behind the FortiGate dialup client and the private network behind this FortiGate dialup server. Because communication cannot be initiated in the opposite direction, there is only one policy.

- Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
- Enter these settings in particular:

Name	Enter an appropriate name for the policy.
Incoming Interface	Select the VPN tunnel (IPsec interface) created in Step 1.
Outgoing Interface	Select the interface that connects to the private network behind this FortiGate unit.
Source	Select all .
Destination Address	Select all .
Action	Select ACCEPT .
NAT	Disable NAT .

Policy-based VPN security policy

- Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
- Enter these settings in particular:

Name	Enter an appropriate name for the policy.
Incoming Interface	Select the interface that connects to the private network behind this FortiGate unit.
Outgoing Interface	Select the FortiGate unit's public interface.
Source	Select the address name that you defined for the private network behind this FortiGate unit.

Destination Address	Select the address name that you defined.
Action	Select IPsec . Under VPN Tunnel , select the name of the Phase 1 configuration that you created in Step " Configuration overview " on page 1745 from the drop-down list. Select Allow traffic to be initiated from the remote site .

- To prevent traffic from the local network from initiating the tunnel after the tunnel has been established, you need to disable the outbound VPN traffic in the CLI

```
config firewall policy
  edit <policy_number>
    set outbound disable
end
```

Place the policy in the policy list above any other policies having similar source and destination addresses.

If configuring a route-based policy, configure a default route for VPN traffic on this interface.

Configuring the FortiGate dialup client

At the FortiGate dialup client, define the Phase 1 parameters needed to authenticate the dialup server and establish a secure connection. See [Phase 1 parameters on page 1655](#).

- Go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
- Edit **Network** (full configuration options are only available once you click the **Convert To Custom Tunnel** button).
- Enter these settings in particular:

Remote Gateway	Select Static IP Address .
IP Address	Type the IP address of the dialup server's public interface.
Interface	Select the interface that connects to the public network.
Mode	The FortiGate dialup client has a dynamic IP address, select Aggressive .
Advanced	Select to view the following options.
Local ID	If you defined a peer ID for the dialup client in the FortiGate dialup server configuration, enter the identifier of the dialup client. The value must be identical to the peer ID that you specified previously in the FortiGate dialup server configuration.

- Edit **Authentication** and enter the following information:

Mode	The FortiGate dialup client has a dynamic IP address, select Aggressive .
-------------	--

- Edit **Phase 1 Proposal** and enter the following information:

Local ID	If you defined a peer ID for the dialup client in the FortiGate dialup server configuration, enter the identifier of the dialup client. The value must be identical to the peer ID that you specified previously in the FortiGate dialup server configuration.
-----------------	--

- Define the Phase 2 parameters needed to create a VPN tunnel with the dialup server. See [Phase 2 parameters on page 1675](#). Enter these settings in particular:

Name	Enter a name to identify this Phase 2 configuration.
Phase 1	Select the name of the Phase 1 configuration that you defined.

- Define names for the addresses or address ranges of the private networks that the VPN links. See [Defining policy addresses on page 1](#). Enter these settings in particular:
 - Define an address name for the server, host, or network behind the FortiGate dialup server.
 - Define an address name for the private network behind the FortiGate dialup client.
- Define security policies to permit communication between the private networks through the VPN tunnel. Route-based and policy-based VPNs require different security policies. For detailed information about creating security policies, see [Defining VPN security policies on page 1](#).

Route-based VPN security policy

Define an ACCEPT security policy to permit communications between hosts on the private network behind this FortiGate dialup client and the private network behind the FortiGate dialup server. Because communication cannot be initiated in the opposite direction, there is only one policy.

- Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
- Enter these settings in particular:

Name	Enter an appropriate name for the policy.
Incoming Interface	Select the interface that connects to the private network behind this FortiGate unit.
Outgoing Interface	Select the VPN tunnel (IPsec interface) created in Step 1.
Source	Select all .
Destination Address	Select all .
Action	Select ACCEPT .
NAT	Disable NAT .

Policy-based VPN security policy

Define an IPsec security policy to permit communications between the source and destination addresses.

- Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
- Enter these settings in particular:

Incoming Interface	Select the interface that connects to the private network behind this FortiGate unit.
Outgoing Interface	Select the FortiGate unit's public interface.
Source	Select the address name that you defined for the private network behind this FortiGate unit.
Destination Address	Select the address name that you defined for the private network behind the dialup server.
Action	<p>Select IPsec. Under VPN Tunnel, select the name of the Phase 1 configuration that you created in Step "Configuration overview " on page 1745 from the drop-down list.</p> <p>Clear Allow traffic to be initiated from the remote site to prevent traffic from the remote network from initiating the tunnel after the tunnel has been established.</p>

Place the policy in the policy list above any other policies having similar source and destination addresses.

IPsec dial-up interface sharing

It is possible to use a single interface for all instances that spawn via a given phase1. In this case, instead of creating an interface per instance, all traffic will run over the single interface and any routes that need creating will be created on that single interface.

The CLI option `"net-device [enable|disable]"` is available in the phase1-interface command sets. Under the new single-interface scheme, instead of relying on routing to guide traffic to the specific instance, all traffic will flow to the specific device and IPsec will need to take care of locating the correct instance for outbound traffic. For this purpose, the CLI option `"tunnel-search"` is provided. The option is only available when the above `"net-device"` option is `"disable"`.

There are two options for `"tunnel-search"`, corresponding to the two ways to select the tunnel for outbound traffic. One is `"selectors"`, meaning selecting a peer using the IPsec selectors (proxy-ids). The other is `"nexthop"` where all the peers use the same default selectors (0/0) while using some routing protocols such as BGP, OSPF, RIPng, etc to resolve the routing. The default for `"tunnel-search"` is `"selectors"`.

Syntax

```

config vpn ipsec phase1-interface
    edit xxx
        set net-device [enable|disable] Enable to create a kernel device for every dialup instance
    next
end
config vpn ipsec phase1-interface
    edit xxx
        set net-device disable
        set tunnel-search [selectors|nexthop] Search for tunnel in selectors or using nexthops
    next
end

```

Supporting IKE Mode Config clients

IKE Mode Config is an alternative to DHCP over IPsec. A FortiGate unit can be configured as either an IKE Mode Config server or client. This chapter contains the following sections:

IKE Mode Config overview

Dialup VPN clients connect to a FortiGate unit that acts as a VPN server, providing the client the necessary configuration information to establish a VPN tunnel. The configuration information typically includes a virtual IP address, netmask, and DNS server address.

IKE Mode Config is available only for VPNs that are route-based, also known as interface-based. A FortiGate unit can function as either an IKE Configuration Method server or client. IKE Mode Config is configurable only in the CLI.

Automatic configuration overview

VPN configuration for remote clients is simpler if it is automated. Several protocols support automatic configuration:

- The Fortinet FortiClient Endpoint Security application can completely configure a VPN connection with a suitably configured FortiGate unit given only the FortiGate unit's address. This protocol is exclusive to Fortinet. For more information, see [FortiClient dialup-client configurations on page 1](#).
- DHCP over IPsec can assign an IP address, Domain, DNS and WINS addresses. The user must first configure IPsec parameters such as gateway address, encryption and authentication algorithms.
- IKE Mode Config can configure host IP address, Domain, DNS and WINS addresses. The user must first configure IPsec parameters such as gateway address, encryption and authentication algorithms. Several network equipment vendors support IKE Mode Config, which is described in the ISAKMP Configuration Method document [draft-dukes-ike-mode-cfg-02.txt](#).

This chapter describes how to configure a FortiGate unit as either an IKE Mode Config server or client.

IKE Mode Config method

IKE Mode Config is configured with the CLI command `config vpn ipsec phase1-interface`. The `mode-cfg` variable enables IKE Mode Config. The `type` field determines whether you are creating an IKE Mode Config server or a client. Setting `type` to `dynamic` creates a server configuration, otherwise the configuration is a client.

Creating an IKE Mode Config client

If the FortiGate unit will connect as a dialup client to a remote gateway that supports IKE Mode Config, the relevant `vpn ipsec phase1-interface` variables are as follows:

Variable	Description
<code>ike-version 1</code>	IKE v1 is the default for FortiGate IPsec VPNs. IKE Mode Config is also compatible with IKE v2 (RFC 4306). Use syntax <code>ike-version 2</code> .
<code>mode-cfg enable</code>	Enable IKE Mode Config.
<code>type {ddns static}</code>	If you set <code>type</code> to <code>dynamic</code> , an IKE Mode Config server is created.
<code>assign-ip {enable disable}</code>	Enable to request an IP address from the server.
<code>interface <interface_name></code>	This is a regular IPsec VPN field. Specify the physical, aggregate, or VLAN interface to which the IPsec tunnel will be bound.
<code>proposal <encryption_combination></code>	This is a regular IPsec VPN field that determines the encryption and authentication settings that the client will accept. For more information, see Phase 1 parameters on page 1655 .
<code>ip-version <4 6></code>	This is a regular IPsec VPN field. By default, IPsec VPNs use IPv4 addressing. You can set <code>ip-version</code> to <code>6</code> to create a VPN with IPv6 addressing.
<code>ipv4-split-exclude</code> <code>ipv6-split-exclude</code>	This command allows the administrator to specify that default traffic flows over the IPsec tunnel except for specified subnets. This is the opposite of the supported <code>split-include</code> feature which allows the administrator to specify that default traffic should not flow over the IPsec tunnel except for specified subnets.

For a complete list of available variables, see the [CLI Reference](#).

IKE Mode Config client example - CLI

In this example, the FortiGate unit connects to a VPN gateway with a static IP address that can be reached through Port 1. Only the port, gateway and proposal information needs to be configured. All other configuration information will come from the IKE Mode Config server.

```
config vpn ipsec phase1-interface
  edit vpn1
    set ip-version 4
    set type static
    set remote-gw <gw_address>
    set interface port 1
    set proposal 3des-sha1 aes128-sha1
    set mode-cfg enable
    set assign-ip enable
  end
```

Creating an IKE Mode Config server

If the FortiGate unit will accept connection requests from dialup clients that support IKE Mode Config, the following `vpn ipsec phase1-interface` settings are required before any other configuration is attempted:

Variable	Description
<code>ike-version 1</code>	IKE v1 is the default for FortiGate IPsec VPNs. IKE Mode Config is also compatible with IKE v2 (RFC 4306). Use syntax <code>ike-version 2</code> .
<code>mode-cfg enable</code>	Enable IKE Mode Config.
<code>type dynamic</code>	Any other setting creates an IKE Mode Config client.
<code>interface <interface_name></code>	This is a regular IPsec VPN field. Specify the physical, aggregate, or VLAN interface to which the IPsec tunnel will be bound.
<code>proposal <encryption_combination></code>	This is a regular IPsec VPN field that determines the encryption and authentication settings that the server will accept. For more information, see Phase 1 parameters on page 1655 .
<code>ip-version <4 6></code>	This is a regular IPsec VPN field. By default, IPsec VPNs use IPv4 addressing. You can set <code>ip-version</code> to <code>6</code> to create a VPN with IPv6 addressing.

IKE Mode Config server example - CLI

In this example, the FortiGate unit assigns IKE Mode Config clients addresses in the range of 10.11.101.160 through 10.11.101.180. DNS and WINS server addresses are also provided. The public interface of the FortiGate unit is Port 1.

When IKE Mode-Configuration is enabled, multiple server IPs can be defined in IPsec Phase 1.

The `ipv4-split-include` variable specifies a firewall address that represents the networks to which the clients will have access. This destination IP address information is sent to the clients.

Only the CLI fields required for IKE Mode Config are shown here. For detailed information about these variables, see the FortiGate CLI Reference.

```
config vpn ipsec phase1-interface
edit "vpn-p1"
    set type dynamic
    set interface "wan1"
    set xauthtype auto
    set mode aggressive
    set mode-cfg enable
    set proposal 3des-sha1 aes128-sha1
    set dpd disable
    set dhgrp 2
    set xauthexpire on-rekey
    set authusrgrp "FG-Group1"
```

```
set ipv4-start-ip 10.10.10.10
set ipv4-end-ip 10.10.10.20
set ipv4-dns-server1 1.1.1.1
set ipv4-dns-server2 2.2.2.2
set ipv4-dns-server3 3.3.3.3
set ipv4-wins-server1 4.4.4.4
set ipv4-wins-server2 5.5.5.5
set domain "fgt1c-domain"
set banner "fgt111C-banner"
set backup-gateway "100.100.100.1" "host1.com" "host2"
set ipv4-split-include OfficeLAN
end
```

IP address assignment

After you have enabled the basic configuration, you can configure IP address assignment for clients, as well as DNS and WINS server assignment. Usually you will want to assign IP addresses to clients.

The simplest method to assign IP addresses to clients is to assign addresses from a specific range, similar to a DHCP server.

If your clients are authenticated by a RADIUS server, you can obtain the user's IP address assignment from the Framed-IP-Address attribute. The user must be authenticated using XAuth.

IKE Mode Config can also use a remote DHCP server to assign the client IP addresses. Up to eight addresses can be selected for either IPv4 or IPv6. After the DHCP proxy has been configured, the `assign-ip-from` command is used to assign IP addresses via DHCP.

Assigning IP addresses from an address range - CLI

If your VPN uses IPv4 addresses,

```
config vpn ipsec phase1-interface
edit vpn1
set mode-cfg-ipversion 4
set assign-ip enable
set assign-ip-type ip
set assign-ip-from range
set ipv4-start-ip <range_start>
set ipv4-end-ip <range_end>
set ipv4-netmask <netmask>
end
```

If your VPN uses IPv6 addresses,

```
config vpn ipsec phase1-interface
edit vpn1
set mode-cfg-ipversion 6
set assign-ip enable
set assign-ip-type ip
set assign-ip-from range
set ipv6-start-ip <range_start>
set ipv6-end-ip <range_end>
end
```

Assigning IP addresses from a RADIUS server - CLI

The users must be authenticated by a RADIUS server and assigned to the FortiGate user group <grpname>. Since the IP address will not be static, `type` is set to `dynamic`, and `mode-cfg` is enabled. This is IKE Configuration Method so that compatible clients can configure themselves with settings that the FortiGate unit provides.

```
config vpn ipsec phase1-interface
  edit vpn1
    set type dynamic
    set mode-cfg enable
    set assign-ip enable
    set assign-ip-from usrgrp
    set xauthtype auto
    set authusrgrp <grpname>
  end
```

Assigning IP address from DHCP - CLI

The DHCP proxy must first be enabled for IKE Mode Config to use DHCP to assign the VPN client IP address(es).

```
config system settings
  set dhcp-proxy enable
  set dhcp-server-ip [ipv4 address]
  set dhcp6-server-ip [ipv6-address]
```

(Up to eight server addresses can be configured)

```
end

config vpn ipsec phase1-interface
  edit vpn1
    set mode-cfg enable
    set assign-ip-from dhcp
  next
end
```

Assigning IP address from a named firewall address/group - CLI

```
config vpn ipsec phase1-interface
  edit <name>vpn1
    set type dynamic
    set assign-ip-from name
    set ipv4-name <name>
    set ipv6-name <name>
  next
end
```

Certificate groups

IKE certificate groups consisting of up to four RSA certificates can be used in IKE Phase 1. Since CA and local certificates are global, the IKE daemon loads them once for all VDOMs and indexes them into trees based on subject and public key hash (for CA certificates), or certificate name (for local certificates). Certificates are linked together based on the issuer, and certificate chains are built by traversing these links. This reduces the need to keep multiple copies of certificates that could exist in multiple chains.

IKE certificate groups can be configured through the CLI.

Configuring the IKE local ID - CLI

```
config vpn certificate local
edit <name>
    set ike-localid <string>
    set ike-localid-type {asn1dn | fqdn}
end
```

Split-exclude in IKEv1 mode-cfg

This feature allows the administrator to specify when using IKEv1 Configuration Method that default traffic flows over the IPsec tunnel except for specified subnets. This is the opposite of the supported `split-include` feature which allows the administrator to specify that default traffic should not flow over the IPsec tunnel except for specified subnets.

The `split-include` and `split-exclude` options can both be specified at the same time. Whether a client does the right thing when both are specified depends on the client.

Syntax

```
config vpn ipsec {phase1 | phase1-interface}
edit <name>
    set ike-version 1
    set type dynamic
    set mode-cfg enable
    set ipv4-split-exclude {all | none | address}
    set ipv6-split-exclude {all | none | address}
next
end
```

Internet-browsing configuration

This section explains how to support secure web browsing performed by dialup VPN clients, and/or hosts behind a remote VPN peer. Remote users can access the private network behind the local FortiGate unit and browse the Internet securely. All traffic generated remotely is subject to the security policy that controls traffic on the private network behind the local FortiGate unit.


The following topics are included in this section:

Configuration overview

A VPN provides secure access to a private network behind the FortiGate unit. You can also enable VPN clients to access the Internet securely. The FortiGate unit inspects and processes all traffic between the VPN clients and hosts on the Internet according to the Internet browsing policy. This is accomplished even though the same FortiGate interface is used for both encrypted VPN client traffic and unencrypted Internet traffic.

In the figure below, FortiGate_1 enables secure Internet browsing for FortiClient Endpoint Security users such as Dialup_1 and users on the Site_2 network behind FortiGate_2, which could be a VPN peer or a dialup client.

Example Internet-browsing configuration



You can adapt any of the following configurations to provide secure Internet browsing:

- A gateway-to-gateway configuration (see [Gateway-to-gateway configurations on page 1](#))
- A FortiClient dialup-client configuration (see [FortiClient dialup-client configurations on page 1](#))
- A FortiGate dialup-client configuration (see [FortiGate dialup-client configurations on page 1](#))

The procedures in this section assume that one of these configurations is in place, and that it is operating properly.

To create an internet-browsing configuration based on an existing gateway-to-gateway configuration, you must edit the gateway-to-gateway configuration as follows:

- On the FortiGate unit that will provide Internet access, create an Internet browsing security policy. See [Configuration overview on page 1759](#), below.
- Configure the remote peer or client to route all traffic through the VPN tunnel. You can do this on a FortiGate unit or on a FortiClient Endpoint Security application. See [Configuration overview on page 1759](#).

Creating an Internet browsing security policy

On the FortiGate unit that acts as a VPN server and will provide secure access to the Internet, you must create an Internet browsing security policy. This policy differs depending on whether your gateway-to-gateway configuration is policy-based or route-based.

Creating an Internet browsing policy - policy-based VPN

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information and then select **OK**:

Name	Enter an appropriate name for the policy.
Incoming Interface	The interface to which the VPN tunnel is bound.
Outgoing Interface	The interface to which the VPN tunnel is bound.
Source	The internal range address of the remote spoke site.
Destination Address	all
Action	Select IPsec . Under VPN Tunnel , select the tunnel that provides access to the private network behind the FortiGate unit. Select Allow traffic to be initiated from the remote site .
NAT	Enable NAT .

Creating an Internet browsing policy - route-based VPN

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information and then select **OK**:

Name	Enter an appropriate name for the policy.
Incoming Interface	The IPsec VPN interface.

Outgoing Interface	The interface that connects to the Internet. The virtual IPsec interface is configured on this physical interface.
Source	The internal range address of the remote spoke site.
Destination Address	all
Action	ACCEPT
NAT	Enable NAT .

The VPN clients must be configured to route all Internet traffic through the VPN tunnel.

Routing all remote traffic through the VPN tunnel

To make use of the Internet browsing configuration on the VPN server, the VPN peer or client must route all traffic through the VPN tunnel. Usually, only the traffic destined for the private network behind the FortiGate VPN server is sent through the tunnel.

The remote end of the VPN can be a FortiGate unit that acts as a peer in a gateway-to-gateway configuration, or a FortiClient application that protects an individual client PC.

- To configure a remote peer FortiGate unit for Internet browsing via VPN, see [Configuring a FortiGate remote peer to support Internet browsing on page 1761](#).
- To configure a FortiClient Endpoint Security application for Internet browsing via VPN, see [Configuring a FortiClient application to support Internet browsing on page 1762](#).

These procedures assume that your VPN connection to the protected private network is working and that you have configured the FortiGate VPN server for Internet browsing as described in [Configuration overview on page 1759](#).

Configuring a FortiGate remote peer to support Internet browsing

The configuration changes to send all traffic through the VPN differ for policy-based and route-based VPNs.

Routing all traffic through a policy-based VPN

1. At the FortiGate dialup client, go to **Policy & Objects > IPv4 Policy**.
2. Select the IPsec security policy and then select **Edit**.
3. From the **Destination Address** list, select **all**.
4. Select **OK**.

Packets are routed through the VPN tunnel, not just those destined for the protected private network.

Routing all traffic through a route-based VPN

1. At the FortiGate dialup client, go to **Network > Static Routes**.
2. Select the default route (destination IP 0.0.0.0) and then select **Edit**. If there is no default route, select **Create New**. Enter the following information and select **OK**:

Destination IP/Mask	Set to Subnet and enter 0.0.0.0/0.0.0.0 in the field provided.
----------------------------	--

Device	Select the IPsec virtual interface.
Administrative Distance	Leave at default.

All packets are routed through the VPN tunnel, not just packets destined for the protected private network.

Configuring a FortiClient application to support Internet browsing

By default, the FortiClient application configures the PC so that traffic destined for the remote protected network passes through the VPN tunnel but all other traffic is sent to the default gateway. You need to modify the FortiClient settings so that it configures the PC to route all outbound traffic through the VPN.

Routing all traffic through VPN - FortiClient application

1. At the remote host, start FortiClient.
2. Go to **Remote Access**.
3. Select the definition that connects FortiClient to the FortiGate dialup server, select the **Settings** icon, and select **Edit the selected connection**.
4. In the **Edit VPN Connection** dialog box, select **Advanced Settings**.
5. In the **Remote Network** group, select **Add**.
6. In the **IP** and **Subnet Mask** fields, type `0.0.0.0/0.0.0.0` and select **OK**.

The address is added to the **Remote Network** list. The first destination IP address in the list establishes a VPN tunnel. The second destination address (`0.0.0.0/0.0.0.0` in this case) forces all other traffic through the VPN tunnel.

7. Select **OK**.

Redundant VPN configurations

This section discusses the options for supporting redundant and partially redundant IPsec VPNs, using route-based approaches.

The following topics are included in this section:

Configuration overview

A FortiGate unit with two interfaces connected to the Internet can be configured to support redundant VPNs to the same remote peer. If the primary connection fails, the FortiGate unit can establish a VPN using the other connection.

Redundant tunnels do not support Tunnel Mode or manual keys. You must use Interface Mode.

A fully-redundant configuration requires redundant connections to the Internet on both peers. The figure below shows an example of this. This is useful to create a reliable connection between two FortiGate units with static IP addresses.


When only one peer has redundant connections, the configuration is partially-redundant. For an example of this, see [Configuration overview on page 1763](#). This is useful to provide reliable service from a FortiGate unit with static IP addresses that accepts connections from dialup IPsec VPN clients.

In a fully-redundant VPN configuration with two interfaces on each peer, four distinct paths are possible for VPN traffic from end to end. Each interface on a peer can communicate with both interfaces on the other peer. This ensures that a VPN will be available as long as each peer has one working connection to the Internet.

You configure a VPN and an entry in the routing table for each of the four paths. All of these VPNs are ready to carry data. You set different routing distances for each route and only the shortest distance route is used. If this route fails, the route with the next shortest distance is used.

The redundant configurations described in this chapter use route-based VPNs, otherwise known as virtual IPsec interfaces. This means that the FortiGate unit must operate in NAT mode. You must use auto-keying. A VPN that is created using manual keys cannot be included in a redundant-tunnel configuration.

The configuration described here assumes that your redundant VPNs are essentially equal in cost and capability. When the original VPN returns to service, traffic continues to use the replacement VPN until the replacement VPN fails. If your redundant VPN uses more expensive facilities, you want to use it only as a backup while the main VPN is down. For information on how to do this, see [Configuration overview on page 1763](#).

Example redundant-tunnel configuration

A VPN that is created using manual keys cannot be included in a redundant-tunnel configuration.

General configuration steps

A redundant configuration at each VPN peer includes:

- One Phase 1 configuration (virtual IPsec interface) for each path between the two peers. In a fully-meshed redundant configuration, each network interface on one peer can communicate with each network interface on the remote peer. If both peers have two public interfaces, this means that each peer has four paths, for example.
- One Phase 2 definition for each Phase 1 configuration.
- One static route for each IPsec interface, with different distance values to prioritize the routes.
- Two Accept security policies per IPsec interface, one for each direction of traffic.
- Dead peer detection enabled in each Phase 1 definition.

The procedures in this section assume that two separate interfaces to the Internet are available on each VPN peer.

Configuring the VPN peers - route-based VPN

VPN peers are configured using Interface Mode for redundant tunnels.

Configure each VPN peer as follows:

1. Ensure that the interfaces used in the VPN have static IP addresses.
2. Create a Phase 1 configuration for each of the paths between the peers.
3. Enable dead peer detection so that one of the other paths is activated if this path fails.
4. Enter these settings in particular, and any other VPN settings as required:

Path 1

Remote Gateway	Select Static IP Address .
IP Address	Type the IP address of the primary interface of the remote peer.
Local Interface	Select the primary public interface of this peer.
Dead Peer Detection	Enable

Path 2

Remote Gateway	Select Static IP Address .
IP Address	Type the IP address of the secondary interface of the remote peer.
Local Interface	Select the primary public interface of this peer.
Dead Peer Detection	Enable

Path 3

Remote Gateway	Select Static IP Address .
IP Address	Type the IP address of the primary interface of the remote peer.
Local Interface	Select the secondary public interface of this peer.
Dead Peer Detection	Enable

Path 4

Remote Gateway	Select Static IP Address .
IP Address	Type the IP address of the secondary interface of the remote peer.
Local Interface	Select the secondary public interface of this peer.
Dead Peer Detection	Enable

For more information, see [Phase 1 parameters on page 1655](#).

5. Create a Phase 2 definition for each path. See [Phase 2 parameters on page 1675](#). Select the Phase 1 configuration (virtual IPsec interface) that you defined for this path. You can select the name from the Static IP Address part of the list.
6. Create a route for each path to the other peer. If there are two ports on each peer, there are four possible paths between the peer devices.

Destination IP/Mask	The IP address and netmask of the private network behind the remote peer.
Device	One of the virtual IPsec interfaces on the local peer.
Distance	For each path, enter a different value to prioritize the paths.

7. Define the security policy for the local primary interface. See [Defining VPN security policies on page 1](#). You need to create two policies for each path to enable communication in both directions. Enter these settings in particular:

Incoming Interface	Select the local interface to the internal (private) network.
Source Address	All
Outgoing Interface	Select one of the virtual IPsec interfaces you created in Step 2.
Destination Address	All
Schedule	Always
Service	Any
Action	ACCEPT

8. Select **Create New**, leave the **Policy Type** as **Firewall** and leave the **Policy Subtype** as **Address**, and enter these settings:

Incoming Interface	Select one of the virtual IPsec interfaces you created in Step 2.
Source Address	All
Outgoing Interface	Select the local interface to the internal (private) network.
Destination Address	All
Schedule	Always
Service	Any
Action	ACCEPT

9. Place the policy in the policy list above any other policies having similar source and destination addresses.
10. Repeat this procedure at the remote FortiGate unit.

Creating a backup IPsec interface

You can configure a route-based VPN that acts as a backup facility to another VPN. It is used only while your main VPN is out of service. This is desirable when the redundant VPN uses a more expensive facility.

You can configure a backup IPsec interface only in the CLI. The backup feature works only on interfaces with static addresses that have dead peer detection enabled. The `monitor` option creates a backup VPN for the specified Phase 1 configuration.

In the following example, `backup_vpn` is a backup for `main_vpn`.

```
config vpn ipsec phase1-interface
  edit main_vpn
    set dpd on
    set interface port1
    set nattraversal enable
    set psksecret "hard-to-guess"
    set remote-gw 192.168.10.8
    set type static
  end
  edit backup_vpn
    set dpd on
    set interface port2
    set monitor main_vpn
    set nattraversal enable
    set psksecret "hard-to-guess"
    set remote-gw 192.168.10.8
    set type static
  end
```

IPsec VPN tunnel aggregate interfaces

This feature allows per-packet routing decisions to be made over two or more IPsec tunnel interfaces, which is usually configured to allow WAN connections to terminate at a data center so that redundancy and load-sharing can be built into this new interface.

The new virtual interface can bond/aggregate IPsec devices and have the new device do round-robin distribution, among other algorithms.

Syntax

```
config vpn ipsec phase1-interface
  edit <name>
    set interface wan1
    set gateway ...
    ...
  next
  edit <name>
    set interface wan2
    set gateway ...
    ...
  next
end
config vpn ipsec phase2-interface
  ...
end
config system interface
```

```
edit ipsec-bond
    set type tun-agg
    set member isp1 isp2
next
end
config router static
edit <value>
    set dst <address>
    set device ipsec-bond
next
end
```

Transparent mode VPNs


This section describes transparent VPN configurations, in which two FortiGate units create a VPN tunnel between two separate private networks transparently.

The following topics are included in this section:

Configuration overview

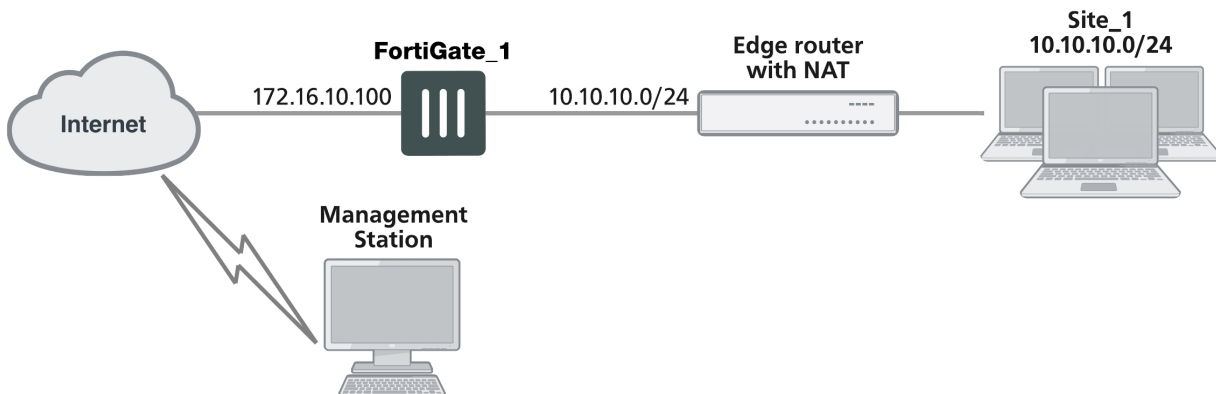
In transparent mode, all interfaces of the FortiGate unit except the management interface (which by default is assigned IP address 10.10.10.1/255.255.255.0) are invisible at the network layer. Typically, when a FortiGate unit runs in transparent mode, different network segments are connected to the FortiGate interfaces. The figure below shows the management station on the same subnet. The management station can connect to the FortiGate unit directly through the web-based manager.

Management station on internal network



An edge router typically provides a public connection to the Internet and one interface of the FortiGate unit is connected to the router. If the FortiGate unit is managed from an external address (see the figure below), the router must translate (NAT) a routable address to direct management traffic to the FortiGate management interface.

Management station on external network



In a transparent VPN configuration, two FortiGate units create a VPN tunnel between two separate private networks transparently. All traffic between the two networks is encrypted and protected by FortiGate security policies.

Both FortiGate units may be running in transparent mode, or one could be running in transparent mode and the other running in NAT mode. If the remote peer is running in NAT mode, it must have a static public IP address.




VPNs between two FortiGate units running in transparent mode do not support inbound/outbound NAT (supported through CLI commands) within the tunnel. In addition, a FortiGate unit running in transparent mode cannot be used in a hub-and-spoke configuration.

Encrypted packets from the remote VPN peer are addressed to the management interface of the local FortiGate unit. If the local FortiGate unit can reach the VPN peer locally, a static route to the VPN peer must be added to the routing table on the local FortiGate unit. If the VPN peer connects through the Internet, encrypted packets from the local FortiGate unit must be routed to the edge router instead. For information about how to add a static route to the FortiGate routing table, see the Advanced Routing Guide.

In the example configuration shown above, Network Address Translation (NAT) is enabled on the router. When an encrypted packet from the remote VPN peer arrives at the router through the Internet, the router performs inbound NAT and forwards the packet to the FortiGate unit. Refer to the software supplier's documentation to configure the router.

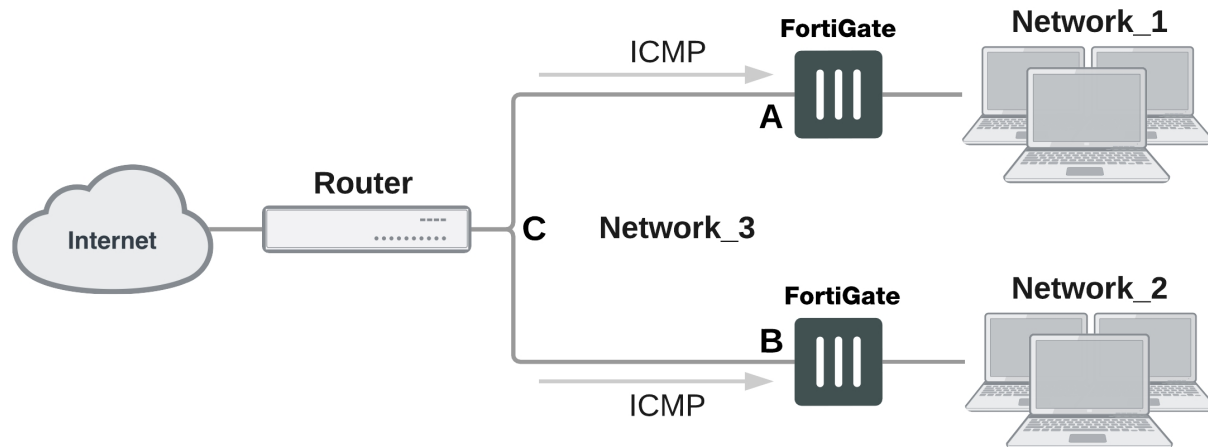
If you want to configure a VPN between two FortiGate units running in transparent mode, each unit must have an independent connection to a router that acts as a gateway to the Internet, and both units must be on separate networks that have a different address space. When the two networks linked by the VPN tunnel have different address spaces (see the figure below), at least one router must separate the two FortiGate units, unless the packets can be redirected using ICMP (as shown in the following figure).

Link between two FortiGate units in transparent mode



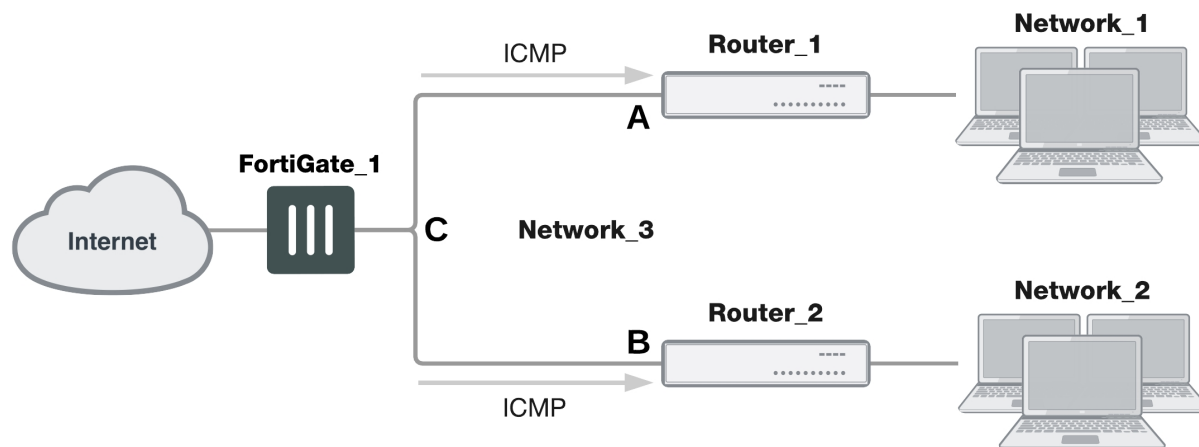
In the figure below, interface C behind the router is the default gateway for both FortiGate units. Packets that cannot be delivered on Network_1 are routed to interface C by default. Similarly, packets that cannot be delivered on Network_2 are routed to interface C. In this case, the router must be configured to redirect packets destined for Network_1 to interface A and redirect packets destined for Network_2 to interface B.

ICMP redirecting packets to two FortiGate units in transparent mode



If there are additional routers behind the FortiGate unit (see the figure below) and the destination IP address of an inbound packet is on a network behind one of those routers, the FortiGate routing table must include routes to those networks. For example, in the following figure, the FortiGate unit must be configured with static routes to interfaces A and B in order to forward packets to Network_1 and Network_2 respectively.

Destinations on remote networks behind internal routers



Transparent VPN infrastructure requirements

- The local FortiGate unit must be operating in transparent mode.
- The management IP address of the local FortiGate unit specifies the local VPN gateway. The management IP address is considered a static IP address for the local VPN peer.
- If the local FortiGate unit is managed through the Internet, or if the VPN peer connects through the Internet, the edge router must be configured to perform inbound NAT and forward management traffic and/or encrypted packets to the FortiGate unit.
- If the remote peer is operating in NAT mode, it must have a static public IP address.

A FortiGate unit operating in transparent mode requires the following basic configuration to operate as a node on the IP network:

- The unit must have sufficient routing information to reach the management station.
- For any traffic to reach external destinations, a default static route to an edge router that forwards packets to the Internet must be present in the FortiGate routing table.
- When all of the destinations are located on the external network, the FortiGate unit may route packets using a single default static route. If the network topology is more complex, one or more static routes in addition to the default static route may be required in the FortiGate routing table.

Only policy-based VPN configurations are possible in transparent mode.

Before you begin

An IPsec VPN definition links a gateway with a tunnel and an IPsec policy. If your network topology includes more than one virtual domain, you must choose components that were created in the same virtual domain. Therefore, before you define a transparent VPN configuration, choose an appropriate virtual domain in which to create the required interfaces, security policies, and VPN components. For more information, see the [Virtual Domains](#) guide.

Configuring the VPN peers

1. The local VPN peer need to operate in transparent mode.
To determine if your FortiGate unit is in transparent mode, go to the **Dashboard > System Information** widget. Select **[change]**. Select transparent for the **Operation Mode**. Two new fields will appear to enter the **Management IP/Netmask**, and the **Default Gateway**.
In transparent mode, the FortiGate unit is invisible to the network. All of its interfaces are on the same subnet and share the same IP address. You only have to configure a management IP address so that you can make configuration changes.

The remote VPN peer may operate in NAT mode or transparent mode.

2. At the local FortiGate unit, define the Phase 1 parameters needed to establish a secure connection with the remote peer. See [Phase 1 parameters on page 1655](#). Select **Advanced** and enter these settings in particular:

Remote Gateway	Select Static IP Address .
IP Address	Type the IP address of the public interface to the remote peer. If the remote peer is a FortiGate unit running in transparent mode, type the IP address of the remote management interface.
Advanced	Select Nat-traversal , and type a value into the Keepalive Frequency field. These settings protect the headers of encrypted packets from being altered by external NAT devices and ensure that NAT address mappings do not change while the VPN tunnel is open. For more information, see Phase 1 parameters on page 1655 and Phase 1 parameters on page 1655 .

3. Define the Phase 2 parameters needed to create a VPN tunnel with the remote peer. See [Phase 2 parameters on page 1675](#). Select the set of Phase 1 parameters that you defined for the remote peer. The name of the remote peer can be selected from the **Static IP Address** list.
4. Define the source and destination addresses of the IP packets that are to be transported through the VPN tunnel. See [Defining VPN security policies on page 1](#). Enter these settings in particular:

- For the originating address (source address), enter the IP address and netmask of the private network behind the local peer network. For the management interface, for example, 10.10.10.0/24. This address needs to be a range to allow traffic from your network through the tunnel. Optionally select `any` for this address.
 - For the remote address (destination address), enter the IP address and netmask of the private network behind the remote peer (for example, 192.168.10.0/24). If the remote peer is a FortiGate unit running in transparent mode, enter the IP address of the remote management interface instead.
5. Define an IPsec security policy to permit communications between the source and destination addresses. See [Defining VPN security policies on page 1](#). Enter these settings in particular:

Incoming Interface	Select the local interface to the internal (private) network.
Source Address	Select the source address that you defined in Step 4.
Outgoing Interface	Select the interface to the edge router. When you configure the IPsec security policy on a remote peer that operates in NAT mode, you select the public interface to the external (public) network instead.
Destination Address	Select the destination address that you defined in Step 4.
VPN Tunnel	<p>Select Use Existing and select the name of the Phase 2 tunnel configuration that you created in Step 3 from the drop-down list.</p> <p>Select Allow traffic to be initiated from the remote site to enable traffic from the remote network to initiate the tunnel.</p>

6. Place the policy in the policy list above any other policies having similar source and destination addresses.
7. Define another IPsec security policy to permit communications between the source and destination addresses in the opposite direction. This security policy and the previous one form a bi-directional policy pair. See [Defining VPN security policies on page 1](#). Enter these settings in particular:

Incoming Interface	Select the interface to the edge router. When you configure the IPsec security policy on a remote peer that operates in NAT mode, you select the public interface to the external (public) network instead.
Source Address	Select the destination address that you defined in Step 4..
Outgoing Interface	Select the local interface to the internal (private) network.
Destination Address	Select the source address that you defined in Step 4.
VPN Tunnel	<p>Select Use Existing and select the name of the Phase 2 tunnel configuration that you created in Step 3 from the drop-down list.</p> <p>Select Allow traffic to be initiated from the remote site to enable traffic from the remote network to initiate the tunnel.</p>

8. Repeat this procedure at the remote FortiGate unit to create bidirectional security policies. Use the local interface and address information local to the remote FortiGate unit.

For more information on transparent mode, see the [System Administration Guide](#).

IPv6 IPsec VPNs

This chapter describes how to configure your FortiGate unit's IPv6 IPsec VPN functionality.



By default IPv6 configurations do not appear on the Web-based Manager. You need to enable the feature first.

To enable IPv6

1. Go to **System > Feature Visibility**.
2. Enable **IPv6**.
3. Select **Apply**.

The following topics are included in this section:

IPv6 IPsec support

FortiOS supports route-based IPv6 IPsec, but not policy-based. This section describes how IPv6 IPsec support differs from IPv4 IPsec support. FortiOS 4.0 MR3 is IPv6 Ready Logo Program Phase 2 certified.

Where both the gateways and the protected networks use IPv6 addresses, sometimes called IPv6 over IPv6, you can create either an auto-keyed or manually-keyed VPN. You can combine IPv6 and IPv4 addressing in an auto-keyed VPN in the following ways:

IPv4 over IPv6	The VPN gateways have IPv6 addresses. The protected networks have IPv4 addresses. The Phase 2 configurations at either end use IPv4 selectors.
IPv6 over IPv4	The VPN gateways have IPv4 addresses. The protected networks use IPv6 addresses. The Phase 2 configurations at either end use IPv6 selectors.

Compared with IPv4 IPsec VPN functionality, there are some limitations:

- Except for IPv6 over IPv4, remote gateways with Dynamic DNS are not supported.
- Selectors cannot be firewall address names. Only IP address, address range and subnet are supported.
- Redundant IPv6 tunnels are not supported.

Certificates

On a VPN with IPv6 Phase 1 configuration, you can authenticate using VPN certificates in which the common name (cn) is an IPv6 address. The `cn-type` keyword of the `user peer` command has an option, `ipv6`, to support this.

Configuration examples

This section consists of the following configuration examples:

- Site-to-site IPv6 over IPv6 VPN example
- Site-to-site IPv6 over IPv4 VPN example
- Site-to-site IPv4 over IPv6 VPN example

Site-to-site IPv6 over IPv6 VPN example

In this example, computers on IPv6-addressed private networks communicate securely over public IPv6 infrastructure.

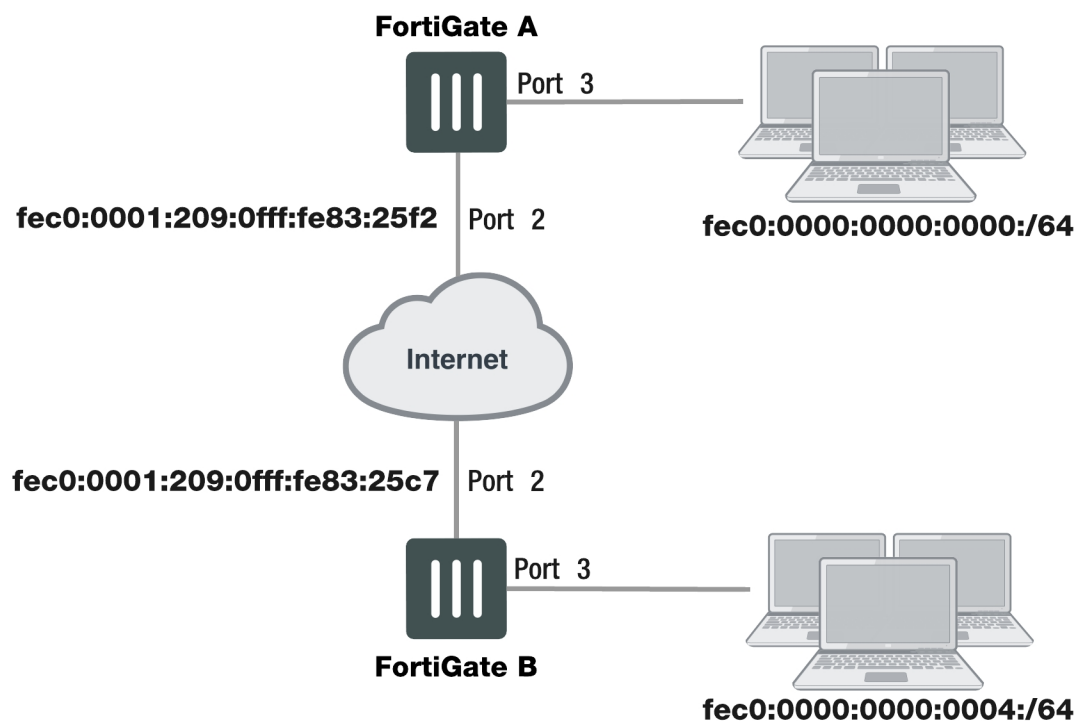


By default IPv6 configurations do not appear on the Web-based Manager. You need to enable the feature first.

To enable IPv6

1. Go to **System > Feature Visibility**.
2. Enable **IPv6**.
3. Select **Apply**.

Example IPv6-over-IPv6 VPN topology



Configure FortiGate A interfaces

Port 2 connects to the public network and port 3 connects to the local network.

```

config system interface
  edit port2
    config ipv6
      set ip6-address fec0::0001:209:0fff:fe83:25f2/64
    end
  next
  edit port3
    config ipv6
      set ip6-address fec0::0000:209:0fff:fe83:25f3/64
    end
  next
end

```

Configure FortiGate A IPsec settings

The Phase 1 configuration creates a virtual IPsec interface on port 2 and sets the remote gateway to the public IP address FortiGate B. This configuration is the same as for an IPv4 route-based VPN, except that `ip-version` is set to 6 and the `remote-gw6` keyword is used to specify an IPv6 remote gateway address.

```

config vpn ipsec phase1-interface
  edit toB
    set ip-version 6
    set interface port2
    set remote-gw6 fec0:0000:0000:0003:209:0fff:fe83:25c7
    set dpd [disable | on-idle | on-demand]
    set psksecret maryhadalittlelamb
    set proposal 3des-md5 3des-sha1
  end

```

By default, Phase 2 selectors are set to accept all subnet addresses for source and destination. The default setting for `src-addr-type` and `dst-addr-type` is `subnet`. The IPv6 equivalent is `subnet6`. The default subnet addresses are 0.0.0.0/0 for IPv4, `:::/0` for IPv6.

```

config vpn ipsec phase2-interface
  edit toB2
    set phase1name toB
    set proposal 3des-md5 3des-sha1
    set pfs enable
    set replay enable
    set src-addr-type subnet6
    set dst-addr-type subnet6
  end

```

Configure FortiGate A security policies

Security policies are required to allow traffic between port3 and the IPsec interface toB in each direction. The address `all6` must be defined using the `firewall address6` command as `::/0`.

```

config firewall policy6
  edit 1
    set srcintf port3
    set dstintf toB
    set srcaddr all6
    set dstaddr all6
    set action accept
    set service ANY
    set schedule always
  end

```

```

next
edit 2
    set srcintf toB
    set dstintf port3
    set srcaddr all6
    set dstaddr all6
    set action accept
    set service ANY
    set schedule always
end

```

Configure FortiGate A routing

This simple example requires just two static routes. Traffic to the protected network behind FortiGate B is routed via the virtual IPsec interface toB. A default route sends all IPv6 traffic out on port2.

```

config router static6
    edit 1
        set device port2
        set dst 0::/0
    next
    edit 2
        set device toB
        set dst fec0:0000:0000:0004::/64
    end
end

```

Configure FortiGate B

The configuration of FortiGate B is very similar to that of FortiGate A. A virtual IPsec interface toA is configured on port2 and its remote gateway is the public IP address of FortiGate A. Security policies enable traffic to pass between the private network and the IPsec interface. Routing ensures traffic for the private network behind FortiGate A goes through the VPN and that all IPv6 packets are routed to the public network.

```

config system interface
    edit port2
        config ipv6
            set ip6-address fec0::0003:209:0fff:fe83:25c7/64
        end
    next
    edit port3
        config ipv6
            set ip6-address fec0::0004:209:0fff:fe83:2569/64
        end
    end
end
config vpn ipsec phase1-interface
    edit toA
        set ip-version 6
        set interface port2
        set remote-gw6 fec0:0000:0000:0001:209:0fff:fe83:25f2
        set dpd [disable | on-idle | on-demand]
        set psksecret maryhadalittlelamb
        set proposal 3des-md5 3des-sha1
    end
end
config vpn ipsec phase2-interface
    edit toA2
        set phase1name toA
        set proposal 3des-md5 3des-sha1
    end
end

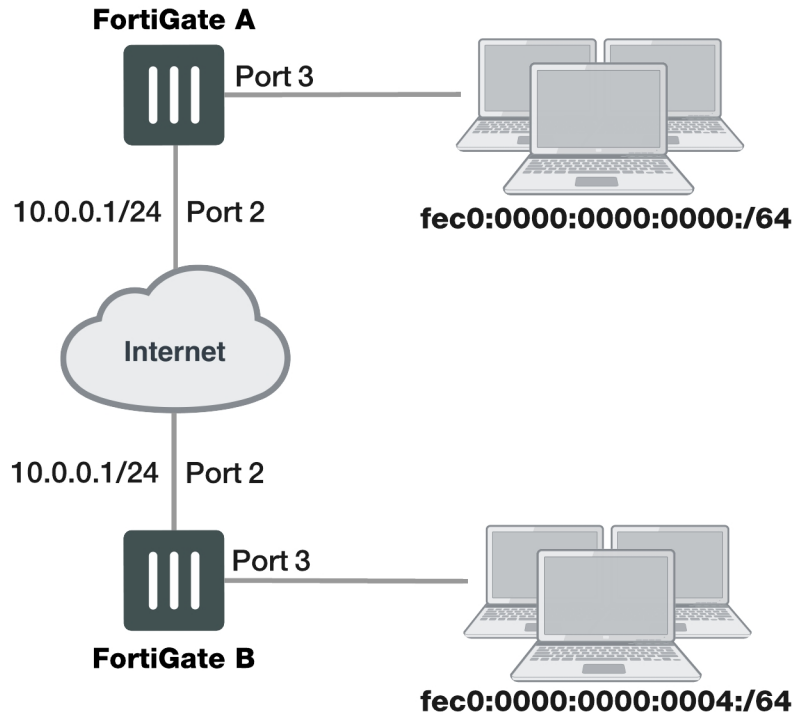
```



```
        set pfs enable
        set replay enable
        set src-addr-type subnet6
        set dst-addr-type subnet6
    end
config firewall policy6
    edit 1
        set srcintf port3
        set dstintf toA
        set srcaddr all6
        set dstaddr all6
        set action accept
        set service ANY
        set schedule always
    next
    edit 2
        set srcintf toA
        set dstintf port3
        set srcaddr all6
        set dstaddr all6
        set action accept
        set service ANY
        set schedule always
    end
config router static6
    edit 1
        set device port2
        set dst 0::/0
    next
    edit 2
        set device toA
        set dst fec0:0000:0000:0000::/64
end
```

Site-to-site IPv6 over IPv4 VPN example

In this example, IPv6-addressed private networks communicate securely over IPv4 public infrastructure.

Example IPv6-over-IPv4 VPN topology**Configure FortiGate A interfaces**

Port 2 connects to the IPv4 public network and port 3 connects to the IPv6 LAN.

```
config system interface
  edit port2
    set 10.0.0.1/24
  next
  edit port3
    config ipv6
      set ip6-address fec0::0001:209:0fff:fe83:25f3/64
    end
  end
```

Configure FortiGate A IPsec settings

The Phase 1 configuration uses IPv4 addressing.

```
config vpn ipsec phase1-interface
  edit toB
    set interface port2
    set remote-gw 10.0.1.1
    set dpd [disable | on-idle | on-demand]
    set psksecret maryhadalittlelamb
    set proposal 3des-md5 3des-sha1
  end
```

The Phase 2 configuration uses IPv6 selectors. By default, Phase 2 selectors are set to accept all subnet addresses for source and destination. The default setting for `src-addr-type` and `dst-addr-type` is `subnet`. The IPv6 equivalent is `subnet6`. The default subnet addresses are `0.0.0.0/0` for IPv4, `::/0` for IPv6.

```
config vpn ipsec phase2-interface
edit toB2
set phase1name toB
set proposal 3des-md5 3des-sha1
set pfs enable
set replay enable
set src-addr-type subnet6
set dst-addr-type subnet6
end
```

Configure FortiGate A security policies

IPv6 security policies are required to allow traffic between `port3` and the IPsec interface `toB` in each direction. Define the address `all6` using the `firewall address6` command as `::/0`.

```
config firewall policy6
edit 1
set srcintf port3
set dstintf toB
set srcaddr all6
set dstaddr all6
set action accept
set service ANY
set schedule always
next
edit 2
set srcintf toB
set dstintf port3
set srcaddr all6
set dstaddr all6
set action accept
set service ANY
set schedule always
end
```

Configure FortiGate A routing

This simple example requires just two static routes. Traffic to the protected network behind FortiGate B is routed via the virtual IPsec interface `toB` using an IPv6 static route. A default route sends all IPv4 traffic, including the IPv4 IPsec packets, out on `port2`.

```
config router static6
edit 1
set device toB
set dst fec0:0000:0000:0004::/64
end
config router static
edit 1
set device port2
set dst 0.0.0.0/0
set gateway 10.0.0.254
end
```

Configure FortiGate B

The configuration of FortiGate B is very similar to that of FortiGate A. A virtual IPsec interface toA is configured on port2 and its remote gateway is the IPv4 public IP address of FortiGate A. The IPsec Phase 2 configuration has IPv6 selectors.

IPv6 security policies enable traffic to pass between the private network and the IPsec interface. An IPv6 static route ensures traffic for the private network behind FortiGate A goes through the VPN and an IPv4 static route ensures that all IPv4 packets are routed to the public network.

```
config system interface
  edit port2
    set 10.0.1.1/24
  next
  edit port3
    config ipv6
      set ip6-address fec0::0004:209:0fff:fe83:2569/64
    end
config vpn ipsec phase1-interface
  edit toA
    set interface port2
    set remote-gw 10.0.0.1
    set dpd [disable | on-idle | on-demand]
    set psksecret maryhadalittlelamb
    set proposal 3des-md5 3des-shal
  end
config vpn ipsec phase2-interface
  edit toA2
    set phase1name toA
    set proposal 3des-md5 3des-shal
    set pfs enable
    set replay enable
    set src-addr-type subnet6
    set dst-addr-type subnet6
  end
config firewall policy6
  edit 1
    set srcintf port3
    set dstintf toA
    set srcaddr all6
    set dstaddr all6
    set action accept
    set service ANY
    set schedule always
  next
  edit 2
    set srcintf toA
    set dstintf port3
    set srcaddr all6
    set dstaddr all6
    set action accept
    set service ANY
    set schedule always
  end
config router static6
  edit 1
```

```

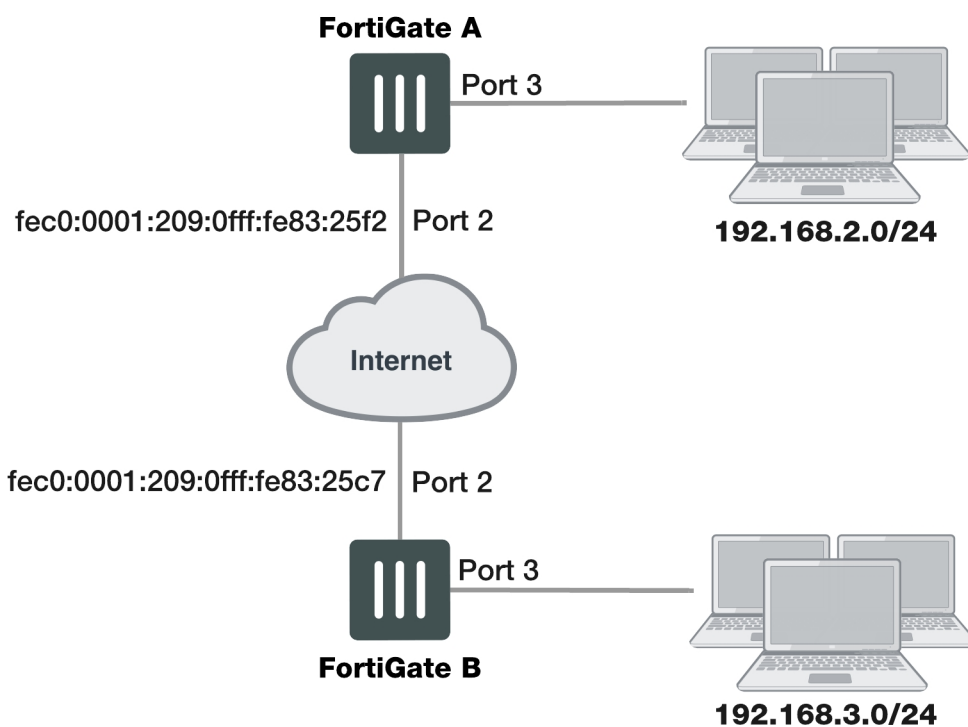
set device toA
set dst fec0:0000:0000:0000::/64
end
config router static
edit 1
set device port2
set gateway 10.0.1.254
end

```

Site-to-site IPv4 over IPv6 VPN example

In this example, two private networks with IPv4 addressing communicate securely over IPv6 infrastructure.

Example IPv4-over-IPv6 VPN topology



Configure FortiGate A interfaces

Port 2 connects to the IPv6 public network and port 3 connects to the IPv4 LAN.

```

config system interface
edit port2
config ipv6
set ip6-address fec0::0001:209:0fff:fe83:25f2/64
end
next
edit port3
set 192.168.2.1/24
end

```

Configure FortiGate A IPsec settings

The Phase 1 configuration is the same as in the IPv6 over IPv6 example.

```
config vpn ipsec phase1-interface
  edit toB
    set ip-version 6
    set interface port2
    set remote-gw6 fec0:0000:0000:0003:209:0fff:fe83:25c7
    set dpd [disable | on-idle | on-demand]
    set psksecret maryhadalittlelamb
    set proposal 3des-md5 3des-sha1
  end
```

The Phase 2 configuration is the same as you would use for an IPv4 VPN. By default, Phase 2 selectors are set to accept all subnet addresses for source and destination.

```
config vpn ipsec phase2-interface
  edit toB2
    set phase1name toB
    set proposal 3des-md5 3des-sha1
    set pfs enable
    set replay enable
  end
```

Configure FortiGate A security policies

Security policies are required to allow traffic between port3 and the IPsec interface toB in each direction. These are IPv4 security policies.

```
config firewall policy
  edit 1
    set srcintf port3
    set dstintf toB
    set srcaddr all
    set dstaddr all
    set action accept
    set service ANY
    set schedule always
  next
  edit 2
    set srcintf toB
    set dstintf port3
    set srcaddr all
    set dstaddr all
    set action accept
    set service ANY
    set schedule always
  end
```

Configure FortiGate A routing

This simple example requires just two static routes. Traffic to the protected network behind FortiGate B is routed via the virtual IPsec interface toB using an IPv4 static route. A default route sends all IPv6 traffic, including the IPv6 IPsec packets, out on port2.

```
config router static6
```

```

edit 1
    set device port2
    set dst 0::/0
next
edit 2
    set device toB
    set dst 192.168.3.0/24
end

```

Configure FortiGate B

The configuration of FortiGate B is very similar to that of FortiGate A. A virtual IPsec interface toA is configured on port2 and its remote gateway is the public IP address of FortiGate A. The IPsec Phase 2 configuration has IPv4 selectors.

IPv4 security policies enable traffic to pass between the private network and the IPsec interface. An IPv4 static route ensures traffic for the private network behind FortiGate A goes through the VPN and an IPv6 static route ensures that all IPv6 packets are routed to the public network.

```

config system interface
    edit port2
        config ipv6
            set ip6-address fec0::0003:fe83:25c7/64
        end
    next
    edit port3
        set 192.168.3.1/24
    end
config vpn ipsec phase1-interface
    edit toA
        set ip-version 6
        set interface port2
        set remote-gw6 fec0:0000:0000:0001:209:0fff:fe83:25f2
        set dpd [disable | on-idle | on-demand]
        set psksecret maryhadalittlelamb
        set proposal 3des-md5 3des-sha1
    end
config vpn ipsec phase2-interface
    edit toA2
        set phaselname toA
        set proposal 3des-md5 3des-sha1
        set pfs enable
        set replay enable
    end
config firewall policy
    edit 1
        set srcintf port3
        set dstintf toA
        set srcaddr all
        set dstaddr all
        set action accept
        set service ANY
        set schedule always
    next
    edit 2
        set srcintf toA
        set dstintf port3

```

```
        set srcaddr all
        set dstaddr all
        set action accept
        set service ANY
        set schedule always
    end
config router static6
    edit 1
        set device port2
        set dst 0::/0
    next
    edit 2
        set device toA
        set dst 192.168.2.0/24
    end
```


L2TP and IPsec (Microsoft VPN)

This section describes how to set up a VPN that is compatible with the Microsoft Windows native VPN, which is Layer 2 Tunneling Protocol (L2TP) with IPsec encryption.

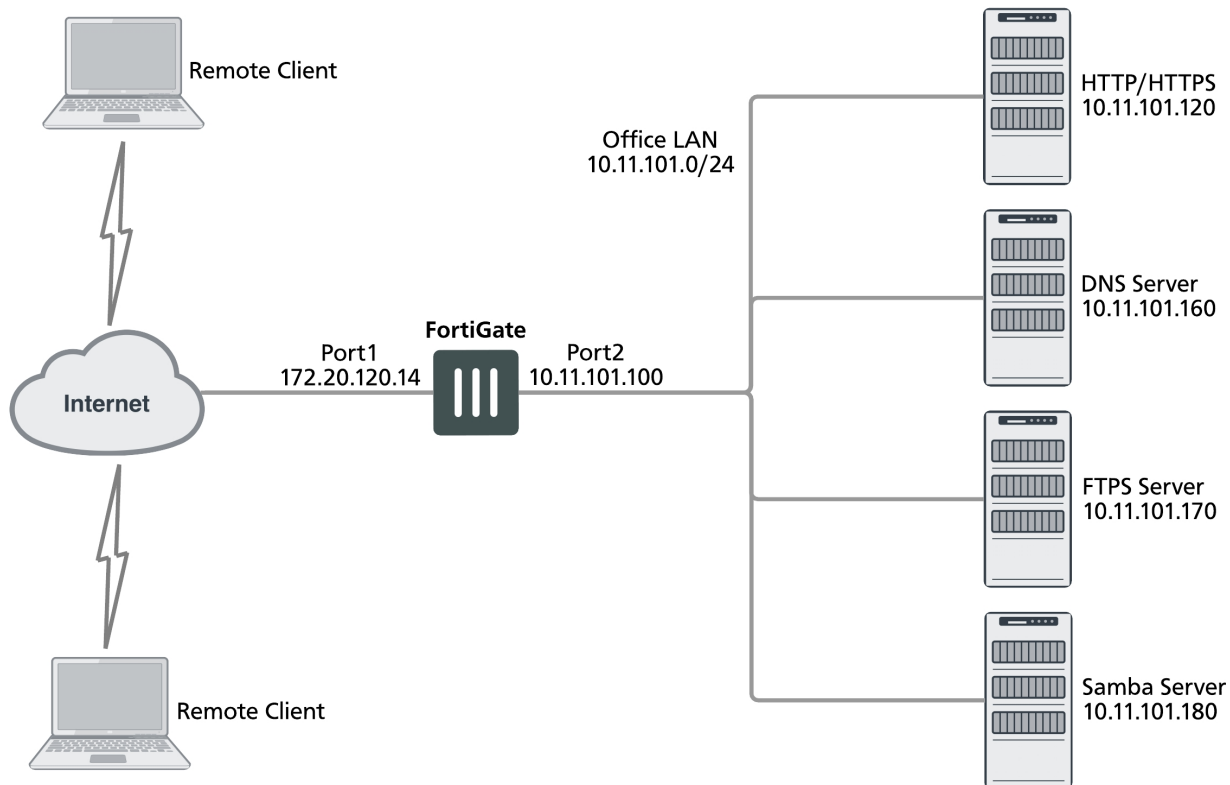
The following topics are included in this section:

For troubleshooting information, refer to [Troubleshooting L2TP and IPsec](#).

Overview

The topology of a VPN for Microsoft Windows dialup clients is very similar to the topology for FortiClient Endpoint Security clients.

Example FortiGate VPN configuration with Microsoft clients



For users, the difference is that instead of installing and using the FortiClient application, they configure a network connection using the software built into the Microsoft Windows operating system. Starting in FortiOS 4.0 MR2, you can configure a FortiGate unit to work with unmodified Microsoft VPN client software.

Layer 2 Tunneling Protocol (L2TP)

L2TP is a tunneling protocol published in 1999 that is used with VPNs, as the name suggests. Microsoft Windows operating system has a built-in L2TP client starting since Windows 2000. Mac OS X 10.3 system and higher also have a built-in client.

L2TP provides no encryption and used UDP port 1701. IPsec is used to secure L2TP packets. The initiator of the L2TP tunnel is called the L2TP Access Concentrator (LAC).

L2TP and IPsec is supported for native Windows XP, Windows Vista and Mac OSX native VPN clients. However, in Mac OSX (OSX 10.6.3, including patch releases) the L2TP feature does not work properly on the Mac OS side.

Assumptions

The following assumptions have been made for this example:

- L2TP protocol traffic is allowed through network firewalls (TCP and UDP port 1701)
- User has Microsoft Windows 2000 or higher — a Windows version that supports L2TP

Configuration overview

The following section consists of configuring the FortiGate unit and configuring the Windows PC.

Configuring the FortiGate unit

To configure the FortiGate unit, you must:

- Configure L2TP users and firewall user group.
- Configure the L2TP VPN, including the IP address range it assigns to clients.
- Configure an IPsec VPN with encryption and authentication settings that match the Microsoft VPN client.
- Configure security policies.

Configuring L2TP users and firewall user group

Remote users must be authenticated before they can request services and/or access network resources through the VPN. The authentication process can use a password defined on the FortiGate unit or an established external authentication mechanism such as RADIUS or LDAP.

Creating user accounts

You need to create user accounts and then add these users to a firewall user group to be used for L2TP authentication. The Microsoft VPN client can automatically send the user's Windows network logon credentials. You might want to use these for their L2TP user name and password.

Creating a user account - web-based manager

1. Go to **User & Device > User Definition** and select **Create New**.
2. Enter the **User Name**.
3. Do one of the following:

- Select **Password** and enter the user's assigned password.
- Select **Match user on LDAP server**, **Match user on RADIUS server**, or **Match user on TACACS+ server** and select the authentication server from the list. The authentication server must be already configured on the FortiGate unit.

4. Select **OK**.

Creating a user account - CLI

To create a user account called `user1` with the password `123_user`, enter:

```
config user local
  edit user1
    set type password
    set passwd "123_user"
    set status enable
  end
```

Creating a user group

When clients connect using the L2TP-over-IPsec VPN, the FortiGate unit checks their credentials against the user group you specify for L2TP authentication. You need to create a firewall user group to use for this purpose.

Creating a user group - web-based manager

1. Go to **User & Device > User Groups**, select **Create New**, and enter the following:

Name	Type or edit the user group name (for example, <code>L2TP_group</code>).
Type	Select Firewall .
Available Users/Groups	The list of Local users, RADIUS servers, LDAP servers, TACACS+ servers, or PKI users that can be added to the user group. To add a member to this list, select the name and then select the right arrow button.
Members	The list of Local users, RADIUS servers, LDAP servers, TACACS+ servers, or PKI users that belong to the user group. To remove a member, select the name and then select the left arrow button.

2. Select **OK**.

Creating a user group - CLI

To create the user group `L2TP_group` and add members `User_1`, `User_2`, and `User_3`, enter:

```
config user group
  edit L2TP_group
    set group-type firewall
    set member User_1 User_2 User_3
  end
```

Configuring L2TP

You can only configure L2TP settings in the CLI. As well as enabling L2TP, you set the range of IP address values that are assigned to L2TP clients and specify the user group that can access the VPN. For example, to allow

access to users in the L2TP_group and assign them addresses in the range 192.168.0.50 to 192.168.0.59, enter:

```
config vpn l2tp
  set sip 192.168.0.50
  set eip 192.168.0.59
  set status enable
  set usrgroup "L2TP_group"
end
```

One of the security policies for the L2TP over IPsec VPN uses the client address range, so you need also need to create a firewall address for that range. For example,

```
config firewall address
  edit L2TPclients
    set type iprange
    set start-ip 192.168.0.50
    set end-ip 192.168.0.59
  end
```

Alternatively, you could define this range in the web-based manager.

Configuring IPsec

The Microsoft VPN client uses IPsec for encryption. The configuration needed on the FortiGate unit is the same as for any other IPsec VPN with the following exceptions.

- Transport mode is used instead of tunnel mode.
- The encryption and authentication proposals must be compatible with the Microsoft client.



Whether Transport mode is *required* depends on the configuration of the peer device (typically an old Windows device, since newer versions of Windows don't require IPsec and L2TP—they can run IPsec natively).



When configuring L2TP, do not name the VPN "L2TP" as that will result in a conflict.

L2TP over IPsec is supported on the FortiGate unit for both policy-based and route-based configurations, but the following example is policy-based.

Configuring Phase 1 - web-based manager

1. Go to **VPN > IPsec Tunnels** and create the new custom tunnel or edit an existing tunnel.
2. Edit the **Phase 1 Proposal** (if it is not available, you may need to click the **Convert to Custom Tunnel** button).

Name	Enter a name for this VPN, dialup_p1 for example.
Remote Gateway	Dialup User
Local Interface	Select the network interface that connects to the Internet. For example, port1.

Mode	Main (ID protection)
Authentication Method	Preshared Key
Pre-shared Key	Enter the preshared key. This key must also be entered in the Microsoft VPN client.
Advanced	Select Advanced to enter the following information.
Phase 1 Proposal	Enter the following Encryption/Authentication pairs: AES256-MD5, 3DES-SHA1, AES192-SHA1
Diffie-Hellman Group	2
NAT Traversal	Enable
Dead Peer Detection	Enable

Configuring Phase 1 - CLI

To create a Phase 1 configuration called dialup_p1 on a FortiGate unit that has port1 connected to the Internet, you would enter:

```
config vpn ipsec phase1
  edit dialup_p1
    set type dynamic
    set interface port1
    set mode main
    set psksecret *****
    set proposal aes256-md5 3des-sha1 aes192-sha1
    set dhgrp 2
    set natTraversal enable
    set dpd [disable | on-idle | on-demand]
  end
```



It is worth noting here that the command `config vpn ipsec phase1` is used rather than `config vpn ipsec phase1-interface` because this configuration is policy-based and not route-based.

Configuring Phase 2 - web-based manager

1. Open the **Phase 2 Selectors** panel.
2. Enter the following information and then select **OK**.

Phase 2 Proposal	Enter the following Encryption/Authentication pairs: AES256-MD5, 3DES-SHA1, AES192-SHA1
Enable replay detection	Enable
Enable perfect forward secrecy (PFS)	Disable

Keylife	3600 seconds
----------------	--------------

3. Make this a transport-mode VPN. You must use the CLI to do this. If your Phase 2 name is dialup_p2, you would enter:

```
config vpn ipsec phase2
  edit dialup_p2
    set encapsulation transport-mode
  end
```

Configuring Phase 2 - CLI

To configure a Phase 2 to work with your phase_1 configuration, you would enter:

```
config vpn ipsec phase2
  edit dialup_p2
    set phase1name dialup_p1
    set proposal aes256-md5 3des-sha1 aes192-sha1
    set replay enable
    set pfs disable
    set keylifeseconds 3600
    set encapsulation transport-mode
  end
```



Once again, note here that the command `config vpn ipsec phase2` is used rather than `config vpn ipsec phase2-interface` because this configuration is policy-based and not route-based.

Configuring security policies

The security policies required for L2TP over IPsec VPN are:

- An IPsec policy, as you would create for any policy-based IPsec VPN
- A regular ACCEPT policy to allow traffic from the L2TP clients to access the protected network

Configuring the IPsec security policy - web-based manager

1. Go to **System > Feature Visibility** and enable **Policy-based IPsec VPN**.
2. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
3. Set the **Action** to **IPsec** and enter the following information:

Incoming Interface	Select the interface that connects to the private network behind this FortiGate unit.
Source Address	All
Outgoing Interface	Select the FortiGate unit's public interface.
Destination Address	All

VPN Tunnel	Select Use Existing and select the name of the Phase 1 configuration that you created. For example, dialup_p1. See Configuring IPsec on page 1789 .
Allow traffic to be initiated from the remote site	enable

4. Select **OK**.

Configuring the IPsec security policy - CLI

If your VPN tunnel (Phase 1) is called dialup_p1, your protected network is on port2, and your public interface is port1, you would enter:

```
config firewall policy
edit 0
set srcintf port2
set dstintf port1
set srcaddr all
set dstaddr all
set action ipsec
set schedule always
set service all
set inbound enable
set vpngateway dialup_p1
end
```

Configuring the ACCEPT security policy - web-based manager

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Leave the **Policy Type** as **Firewall** and leave the **Policy Subtype** as **Address**.
3. Enter the following information and select **OK**:

Incoming Interface	Select the FortiGate unit's public interface.
Source Address	Select the firewall address that you defined for the L2TP clients.
Outgoing Interface	Select the interface that connects to the private network behind this FortiGate unit.
Destination Address	All
Action	ACCEPT

Configuring the ACCEPT security policy - CLI

If your public interface is port1, your protected network is on port2, and L2TPclients is the address range that L2TP clients use, you would enter:

```
config firewall policy
edit 1
set srcintf port1
set dstintf port2
set srcaddr L2TPclients
set dstaddr all
```

```
set action accept
set schedule always
set service all
end
```

Configuring the Windows PC

Configuration of the Windows PC for a VPN connection to the FortiGate unit consists of the following:

1. In Network Connections, configure a Virtual Private Network connection to the FortiGate unit.
2. Ensure that the IPSEC service is running.
3. Ensure that IPsec has not been disabled for the VPN client. It may have been disabled to make the Microsoft VPN compatible with an earlier version of FortiOS.

The instructions in this section are based on Windows XP. Other versions of Windows may vary slightly.

Configuring the network connection

1. Open **Network Connections**.
This is available through the Control Panel.
2. Double-click **New Connection Wizard** and **Select Next**.
3. Select **Connect to the network at my workplace**.
4. Select **Next**.
5. Select **Virtual Private Network connection** and select **Next**.
6. In the **Company Name** field, enter a name for the connection and select **Next**.
7. Select **Do not dial the initial connection** and then select **Next**.
8. Enter the public IP address or FQDN of the FortiGate unit and select **Next**.
9. Optionally, select **Add a shortcut to this connection to my desktop**.
10. Select **Finish**.
The **Connect** dialog opens on the desktop.
11. Select **Properties** and then select the **Security** tab.
12. Select **IPsec Settings**.
13. Select **Use pre-shared key for authentication**, enter the preshared key that you configured for your VPN, and select **OK**.
14. Select **OK**.

Checking that the IPsec service is running

1. Open **Administrative Tools** through the Control Panel.
2. Double-click **Services**.
3. Look for IPSEC Services. Confirm that the **Startup Type** is **Automatic** and **Status** is set to **Started**. If needed, double-click **IPsec Services** to change these settings.

Checking that IPsec has not been disabled

1. Select **Start > Run**.
2. Enter regedit and select **OK**.
3. Find the Registry key HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\RasMan\Parameters
4. If there is a ProhibitIPsec value, it must be set to 0.

Enforcing IPsec in L2TP configuration

An `enforce-ipsec` option is available in L2TP configuration to force the FortiGate L2TP server to accept only IPsec encrypted connections.

Syntax

```
config vpn l2tp
  set eip 50.0.0.100
  set sip 50.0.0.1
  set status enable
  set enforce-ipsec-interface {disable | enable}      (default = disable)
  set usrgrp <group_name>
end
```

GRE over IPsec (Cisco VPN)

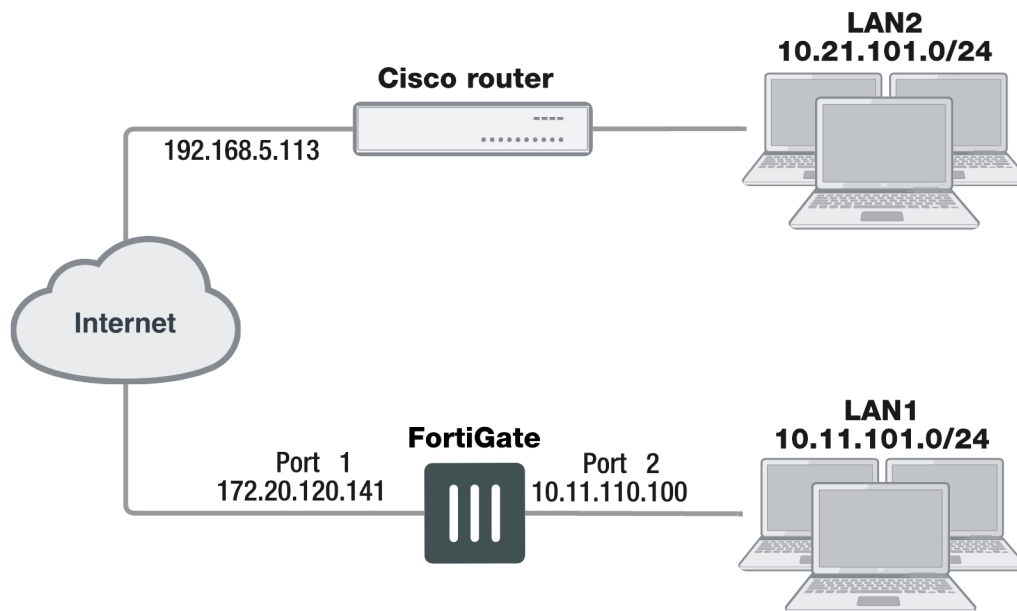
This section describes how to configure a FortiGate VPN that is compatible with Cisco-style VPNs that use GRE in an IPsec tunnel.

The following topics are included in this section:

Cisco products that include VPN support often use Generic Routing Encapsulation (GRE) protocol tunnel over IPsec encryption. This chapter describes how to configure a FortiGate unit to work with this type of Cisco VPN.

Cisco VPNs can use either transport mode or tunnel mode IPsec. Before FortiOS 4.0 MR2, the FortiGate unit was compatible only with tunnel mode IPsec.

Example FortiGate to Cisco GRE-over-IPsec VPN



In this example, users on LAN1 are provided access to LAN2.

Configuration overview

The following section consists of configuring the FortiGate unit and configuring the Cisco router.

Configuring the FortiGate unit

There are several steps to the GRE-over-IPsec configuration:

- Enable overlapping subnets. This is needed because the IPsec and GRE tunnels will use the same addresses.
- Configure a route-based IPsec VPN on the external interface.
- Configure a GRE tunnel on the virtual IPsec interface. Set its local gateway and remote gateway addresses to match the local and remote gateways of the IPsec tunnel.
- Configure security policies to allow traffic to pass in both directions between the GRE virtual interface and the IPsec virtual interface.
- Configure security policies to allow traffic to pass in both directions between the protected network interface and the GRE virtual interface.
- Configure a static route to direct traffic destined for the network behind the Cisco router into the GRE-over-IPsec tunnel.

Enabling overlapping subnets

By default, each FortiGate unit network interface must be on a separate network. The configuration described in this chapter assigns an IPsec tunnel end point and the external interface to the same network. Enable subnet overlap as follows:

```
config system settings
    set allow-subnet-overlap enable
end
```

Configuring the IPsec VPN

A route-based VPN is required. It must use encryption and authentication algorithms compatible with the Cisco equipment to which it connects. In this chapter, preshared key authentication is shown.

Configuring the IPsec VPN - web-based manager

1. Define the Phase 1 configuration needed to establish a secure connection with the remote Cisco device. Enter these settings in particular:

Name	Enter a name to identify the VPN tunnel, tocsico for example. This is the name of the virtual IPsec interface. It appears in Phase 2 configurations, security policies and the VPN monitor.
Remote Gateway	Select Static IP Address .
IP Address	Enter the IP address of the Cisco device public interface. For example, 192.168.5.113.
Local Interface	Select the FortiGate unit's public interface. For example, 172.20.120.141.

Mode	Select Main (ID Protection) .
Authentication Method	Preshared Key
Pre-shared Key	Enter the preshared key. It must match the preshared key on the Cisco device.
Advanced	Select the Advanced button to see the following settings.
Phase 1 Proposal	3DES-MD5
	At least one proposal must match the settings on the Cisco unit.

For more information about these settings, see [Phase 1 parameters on page 1655](#).

2. Define the Phase 2 parameters needed to create a VPN tunnel with the remote peer. For compatibility with the Cisco router, Quick Mode Selectors must be entered, which includes specifying protocol 47, the GRE protocol. Enter these settings in particular:

Phase 2 Proposal	3DES-MD5
	At least one proposal must match the settings on the Cisco unit.
Quick Mode Selector	
Source Address	Enter the GRE local tunnel end IP address. For example 172.20.120.141.
Source Port	0
Destination Address	Enter the GRE remote tunnel end IP address. For example 192.168.5.113.
Destination Port	0
Protocol	47

For more information about these settings, see [Phase 2 parameters on page 1675](#).

3. If the Cisco device is configured to use transport mode IPsec, you need to use transport mode on the FortiGate VPN. You can configure this only in the CLI. In your Phase 2 configuration, set `encapsulation to transport-mode` as follows:

```
config vpn phase2-interface
  edit to_cisco_p2
    set encapsulation transport-mode
  end
```

Configuring the IPsec VPN - CLI

```
config vpn ipsec phase1-interface
  edit tocisco
    set interface port1
    set proposal 3des-sha1 aes128-sha1
```

```

        set remote-gw 192.168.5.113
        set psksecret xxxxxxxxxxxxxxxxx
    end
config vpn ipsec phase2-interface
edit tocisco_p2
    set phase1name "tocisco"
    set proposal 3des-md5
    set encapsulation tunnel-mode // if tunnel mode
    set encapsulation transport-mode // if transport mode
    set protocol 47
    set src-addr-type ip
    set dst-start-ip 192.168.5.113
    set src-start-ip 172.20.120.141
end

```

Adding IPsec tunnel end addresses

The Cisco configuration requires an address for its end of the IPsec tunnel. The addresses are set to match the GRE gateway addresses. Use the CLI to set the addresses, like this:

```

config system interface
edit tocisco
    set ip 172.20.120.141 255.255.255.255
    set remote-ip 192.168.5.113
end

```

Configuring the GRE tunnel

The GRE tunnel runs between the virtual IPsec public interface on the FortiGate unit and the Cisco router. You must use the CLI to configure a GRE tunnel. In the example, you would enter:

```

config system gre-tunnel
edit gre1
    set interface tocisco
    set local-gw 172.20.120.141
    set remote-gw 192.168.5.113
end

```

`interface` is the virtual IPsec interface, `local-gw` is the FortiGate unit public IP address, and `remote-gw` is the remote Cisco device public IP address

Adding GRE tunnel end addresses

You will also need to add tunnel end addresses. The Cisco router configuration requires an address for its end of the GRE tunnel. Using the CLI, enter tunnel end addresses that are not used elsewhere on the FortiGate unit, like this:

```

config system interface
edit gre1
    set ip 10.0.1.1 255.255.255.255
    set remote-ip 10.0.1.2
end

```

Configuring security policies

Two sets of security policies are required:

- Policies to allow traffic to pass in both directions between the GRE virtual interface and the IPsec virtual interface.
- Policies to allow traffic to pass in both directions between the protected network interface and the GRE virtual interface.

Configuring security policies - web-based manager

1. Define an ACCEPT firewall security policy to permit communications between the protected network and the GRE tunnel:

Incoming Interface	Select the interface that connects to the private network behind this FortiGate unit.
Source Address	All
Outgoing Interface	Select the GRE tunnel virtual interface you configured.
Destination Address	All
Action	ACCEPT
Enable NAT	Disable

2. To permit the remote client to initiate communication, you need to define a firewall address security policy for communication in that direction:

Incoming Interface	Select the GRE tunnel virtual interface you configured.
Source Address	All
Outgoing Interface	Select the interface that connects to the private network behind this FortiGate unit.
Destination Address	All
Action	ACCEPT
Enable NAT	Disable

3. Define a pair of ACCEPT firewall address security policies to permit traffic to flow between the GRE virtual interface and the IPsec virtual interface:

Incoming Interface	Select the GRE virtual interface. See Configuring the GRE tunnel on page 1798 .
Source Address	All
Outgoing Interface	Select the virtual IPsec interface you created. See Configuring the IPsec VPN on page 1796 .
Destination Address	All
Action	ACCEPT
Enable NAT	Disable

Incoming Interface	Select the virtual IPsec interface you created. See Configuring the IPsec VPN on page 1796 .
Source Address	All
Outgoing Interface	Select the GRE virtual interface. See Configuring the GRE tunnel on page 1798 .
Destination Address	All
Action	ACCEPT
Enable NAT	Disable

Configuring security policies - CLI

```

config firewall policy
  edit 1 // LAN to GRE tunnel
    set srcintf port2
    set dstintf gre1
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
  next
  edit 2 // GRE tunnel to LAN
    set srcintf gre1
    set dstintf port2
    set srcaddr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
  next
  edit 3 // GRE tunnel to IPsec interface
    set srcintf "gre1"
    set dstintf "tocisco"
    set srcaddr "all"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY"
  next
  edit 4 // IPsec interface to GRE tunnel
    set srcintf "tocisco"
    set dstintf "gre1"
    set srcaddr "all"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY"
end

```

Configuring routing

Traffic destined for the network behind the Cisco router must be routed to the GRE tunnel. To do this, create a static route

1. Go to **Network > Static Routes** and select **Create New**.
2. Enter the following information and select **OK**.

Destination IP/Mask	Enter the IP address and netmask for the network behind the Cisco router. For example 10.21.101.0 255.255.255.0.
Device	Select the GRE virtual interface.
Distance (Advanced)	Leave setting at default value.

In the CLI, using the example values, you would enter

```
config router static
edit 0
set device gre1
set dst 10.21.101.0 255.255.255.0
end
```

Changing GRE over GRE tunnel interface attributes

Administrators can change GRE over GRE tunnel attributes, such as assigning an IP address for a specific configuration application, even if the child interface is not an IPsec tunnel interface.

IPv6 support for GRE tunnels

Support is provided for GRE tunnel termination using IPv6 addresses on both ends of the tunnel (similar to IPv4 functionality).

Syntax

```
config system gre-tunnel
edit <name>
set ip-version 6
set remote-gw6 11:1:1::1
set local-gw6 11:1:1::2
...
next
end
```

Configuring the Cisco router

Using Cisco IOS, you would configure the Cisco router as follows, using the addresses from the example:

```
config ter
crypto ipsec transform-set myset esp-3des esp-md5-hmac
no mode
exit
no ip access-list extended tunnel
ip access-list extended tunnel
```



```
permit gre host 192.168.5.113 host 172.20.120.141
exit
interface Tunnel1
ip address 10.0.1.2 255.255.255.0
tunnel source 192.168.5.113
tunnel destination 172.20.120.141
!
ip route 10.11.101.0 255.255.255.0 Tunnel1
end
clear crypto sa
clear crypto isakmp
```

For transport mode, change `no mode to mode transport`.

This is only the portion of the Cisco router configuration that applies to the GRE-over-IPsec tunnel. For more information, refer to the Cisco documentation.

Keep-alive support for GRE

The FortiGate can send a GRE keep-alive response to a Cisco device to detect a GRE tunnel. If it fails, it will remove any routes over the GRE interface.

Syntax

```
config system gre-tunnel
edit <id>
    set keepalive-interval <value: 0-32767>
    set keepalive-failtimes <value: 1-255>
next
end
```

Protecting OSPF with IPsec

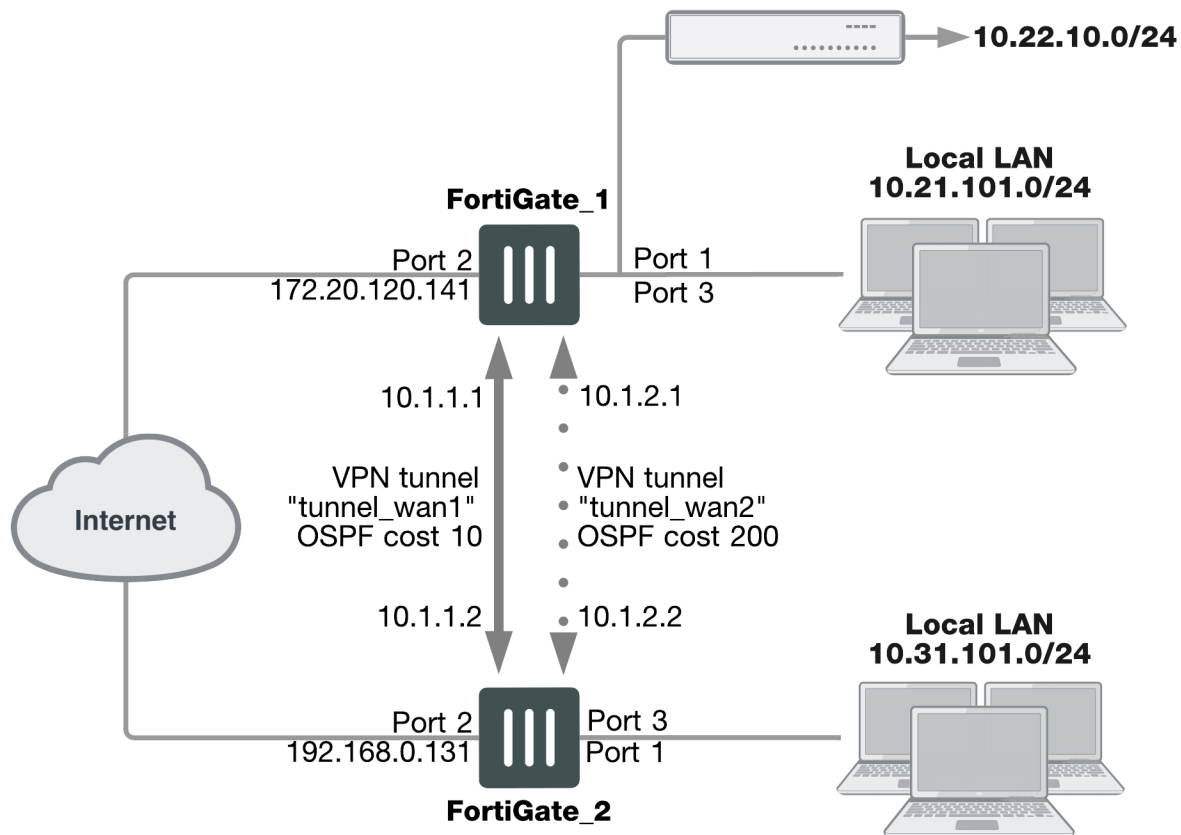
For enhanced security, OSPF dynamic routing can be carried over IPsec VPN links.

The following topics are included in this section:

Configuration overview

This chapter shows an example of OSPF routing conducted over an IPsec tunnel between two FortiGate units. The network shown below is a single OSPF area. FortiGate_1 is an Area border router that advertises a static route to 10.22.10.0/24 in OSPF. FortiGate_2 advertises its local LAN as an OSPF internal route.

OSPF over an IPsec VPN tunnel



The section [Configuration overview](#) describes the configuration with only one IPsec VPN tunnel, tunnel_wan1. Then, the section [Configuration overview](#) describes how you can add a second tunnel to provide a redundant backup path. This is shown above as VPN tunnel "tunnel_wan2".

Only the parts of the configuration concerned with creating the IPsec tunnel and integrating it into the OSPF network are described. It is assumed that security policies are already in place to allow traffic to flow between the interfaces on each FortiGate unit.

OSPF over IPsec configuration

There are several steps to the OSPF-over-IPsec configuration:

- Configure a route-based IPsec VPN on an external interface. It will connect to a corresponding interface on the other FortiGate unit. Define the two tunnel-end addresses.
- Configure a static route to the other FortiGate unit.
- Configure the tunnel network as part of the OSPF network and define the virtual IPsec interface as an OSPF interface.

This section describes the configuration with only one VPN, tunnel_wan1. The other VPN is added in the section [Configuration overview on page 1803](#).

Configuring the IPsec VPN

A route-based VPN is required. In this chapter, preshared key authentication is shown. Certificate authentication is also possible. Both FortiGate units need this configuration.

Configuring Phase 1

1. Define the Phase 1 configuration needed to establish a secure connection with the other FortiGate unit. For more information, see [Phase 1 parameters on page 1655](#).

Enter these settings in particular:

Name	Enter a name to identify the VPN tunnel, tunnel_wan1 for example. This becomes the name of the virtual IPsec interface.
Remote Gateway	Select Static IP Address .
IP Address	Enter the IP address of the other FortiGate unit's public (Port 2) interface.
Local Interface	Select this FortiGate unit's public (Port 2) interface.
Mode	Select Main (ID Protection) .
Authentication Method	Preshared Key
Pre-shared Key	Enter the preshared key. It must match the preshared key on the other FortiGate unit.
Advanced	Select Advanced .

Assigning the tunnel end IP addresses

1. Go to **Network > Interfaces**, select the virtual IPsec interface that you just created on Port 2 and select **Edit**.
2. In the **IP** and **Remote IP** fields, enter the following tunnel end addresses:

	FortiGate_1	FortiGate_2
IP	10.1.1.1	10.1.1.2
Remote_IP	10.1.1.2	10.1.1.1

These addresses are from a network that is not used for anything else.

Configuring Phase 2

1. Enter a name to identify this Phase 2 configuration, `twan1_p2`, for example.
2. Select the name of the Phase 1 configuration that you defined in Step ["Configuration overview" on page 1803](#), `tunnel_wan1` for example.

Configuring static routing

You need to define the route for traffic leaving the external interface.

1. Go to **Network > Static Routes**, select **Create New**.
2. Enter the following information.

Destination IP/Mask	Leave as 0.0.0.0 0.0.0.0.
Device	Select the external interface.
Gateway	Enter the IP address of the next hop router.

Configuring OSPF

This section does not attempt to explain OSPF router configuration. It focusses on the integration of the IPsec tunnel into the OSPF network. This is accomplished by assigning the tunnel as an OSPF interface, creating an OSPF route to the other FortiGate unit.

This configuration uses loopback interfaces to ease OSPF troubleshooting. The OSPF router ID is set to the loopback interface address. The loopback interface ensures the router is always up. Even though technically the router ID doesn't have to match a valid IP address on the FortiGate unit, having an IP that matches the router ID makes troubleshooting a lot easier.

The two FortiGate units have slightly different configurations. FortiGate_1 is an AS border router that advertises its static default route. FortiGate_2 advertises its local LAN as an OSPF internal route.

Setting the router ID for each FortiGate unit to the lowest possible value is useful if you want the FortiGate units to be the designated router (DR) for their respective ASes. This is the router that broadcasts the updates for the AS.

Leaving the IP address on the OSPF interface at 0.0.0.0 indicates that all potential routes will be advertised, and it will not be limited to any specific subnet. For example if this IP address was 10.1.0.0, then only routes that match that subnet will be advertised through this interface in OSPF.

FortiGate_1 OSPF configuration

When configuring FortiGate_1 for OSPF, the loopback interface is created, and then you configure OSPF area networks and interfaces.

With the exception of creating the loopback interface, OSPF for this example can all be configured in either the web-based manager or CLI.

Creating the loopback interface

A loopback interface can be configured in the CLI only. For example, if the interface will have an IP address of 10.0.0.1, you would enter:

```
config system interface
  edit lback1
    set vdom root
    set ip 10.0.0.1 255.255.255.255
    set type loopback
  end
```

The loopback addresses and corresponding router IDs on the two FortiGate units must be different. For example, set the FortiGate 1 loopback to 10.0.0.1 and the FortiGate 2 loopback to 10.0.0.2.

Configuring OSPF area, networks, and interfaces - web-based manager

1. On FortiGate_1, go to **Network > OSPF**.
2. Enter the following information to define the router, area, and interface information.

Router ID	Enter 10.0.0.1. Select Apply before entering the remaining information.
Advanced Options	
Redistribute	Select the Connected and Static check boxes. Use their default metric values.
Areas	Select Create New , enter the Area and Type and then select OK .
Area	0.0.0.0
Type	Regular
Interfaces	Enter a name for the OSPF interface, ospf_wan1 for example.
Name	
Interface	Select the virtual IPsec interface, tunnel_wan1.
IP	0.0.0.0

3. For **Networks**, select **Create New**.
4. Enter the **IP/Netmask** of 10.1.1.0/255.255.255.0 and an **Area** of 0.0.0.0.
5. For **Networks**, select **Create New**.
6. Enter the **IP/Netmask** of 10.0.0.1/255.255.255.0 and an **Area** of 0.0.0.0.
7. Select **Apply**.

Configuring OSPF area and interfaces - CLI

Your loopback interface is 10.0.0.1, your tunnel ends are on the 10.1.1.0/24 network, and your virtual IPsec interface is named `tunnel_wan1`. Enter the following CLI commands:

```

config router ospf
  set router-id 10.0.0.1
  config area
    edit 0.0.0.0
  end
  config network
    edit 4
      set prefix 10.1.1.0 255.255.255.0
    next
    edit 2
      set prefix 10.0.0.1 255.255.255.255
    end
  config ospf-interface
    edit ospf_wan1
      set cost 10
      set interface tunnel_wan1
      set network-type point-to-point
    end
  config redistribute connected
    set status enable
  end
  config redistribute static
    set status enable
  end
end
end

```

FortiGate_2 OSPF configuration

When configuring FortiGate_2 for OSPF, the loopback interface is created, and then you configure OSPF area networks and interfaces.

Configuring FortiGate_2 differs from FortiGate_1 in that three interfaces are defined instead of two. The third interface is the local LAN that will be advertised into OSPF.

With the exception of creating the loopback interface, OSPF for this example can all be configured in either the web-based manager or CLI.

Creating the loopback interface

A loopback interface can be configured in the CLI only. For example, if the interface will have an IP address of 10.0.0.2, you would enter:

```

config system interface
  edit lback1
    set vdom root
    set ip 10.0.0.2 255.255.255.255
    set type loopback
  end

```

The loopback addresses on the two FortiGate units must be different. For example, set the FortiGate 1 loopback to 10.0.0.1 and the FortiGate 2 loopback to 10.0.0.2.

Configuring OSPF area and interfaces - web-based manager

1. On FortiGate_2, go to **Network > OSPF**.
2. Complete the following.

Router ID	10.0.0.2
Areas	Select Create New , enter the Area and Type and then select OK .
Area	0.0.0.0
Type	Regular
Interfaces	
Name	Enter a name for the OSPF interface, ospf_wan1 for example.
Interface	Select the virtual IPsec interface, tunnel_wan1.
IP	0.0.0.0

- For **Networks**, select **Create New**.
- Enter the following information for the loopback interface:

IP/Netmask	10.0.0.2/255.255.255.255
Area	0.0.0.0

- For **Networks**, select **Create New**.
- Enter the following information for the tunnel interface:

IP/Netmask	10.1.1.0/255.255.255.255
Area	0.0.0.0

- For **Networks**, select **Create New**.
- Enter the following information for the local LAN interface:

IP/Netmask	10.31.101.0/255.255.255.255
Area	0.0.0.0

- Select **Apply**.

Configuring OSPF area and interfaces - CLI

If for example, your loopback interface is 10.0.0.2, your tunnel ends are on the 10.1.1.0/24 network, your local LAN is 10.31.101.0/24, and your virtual IPsec interface is named tunnel_wan1, you would enter:

```
config router ospf
  set router-id 10.0.0.2
  config area
    edit 0.0.0.0
  end
  config network
    edit 1
      set prefix 10.1.1.0 255.255.255.0
    next
```

```
edit 2
    set prefix 10.31.101.0 255.255.255.0
next
edit 2
    set prefix 10.0.0.2 255.255.255.255
end
config ospf-interface
    edit ospf_wan1
        set interface tunnel_wan1
        set network-type point-to-point
    end
end
```

Creating a redundant configuration

You can improve the reliability of the OSPF over IPsec configuration described in the previous section by adding a second IPsec tunnel to use if the default one goes down. Redundancy in this case is not controlled by the IPsec VPN configuration but by the OSPF routing protocol.

To do this you:

- Create a second route-based IPsec tunnel on a different interface and define tunnel end addresses for it.
- Add the tunnel network as part of the OSPF network and define the virtual IPsec interface as an additional OSPF interface.
- Set the OSPF cost for the added OSPF interface to be significantly higher than the cost of the default route.

Adding the second IPsec tunnel

The configuration is the same as in [Configuring the IPsec VPN on page 1804](#), but the interface and addresses will be different. Ideally, the network interface you use is connected to a different Internet service provider for added redundancy.

When adding the second tunnel to the OSPF network, choose another unused subnet for the tunnel ends, 10.1.2.1 and 10.1.2.2 for example.

Adding the OSPF interface

OSPF uses the metric called cost when determining the best route, with lower costs being preferred. Up to now in this example, only the default cost of 10 has been used. Cost can be set only in the CLI.

The new IPsec tunnel will have its OSPF cost set higher than that of the default tunnel to ensure that it is only used if the first tunnel goes down. The new tunnel could be set to a cost of 200 compared to the default cost is 10. Such a large difference in cost will ensure this new tunnel will only be used as a last resort.

If the new tunnel is called tunnel_wan2, you would enter the following on both FortiGate units:

```
config router ospf
    config ospf-interface
        edit ospf_wan2
            set cost 200
            set interface tunnel_wan2
            set network-type point-to-point
        end
    end
```


Redundant OSPF routing over IPsec

This example sets up redundant secure communication between two remote networks using an Open Shortest Path First (OSPF) VPN connection. In this example, the HQ FortiGate unit will be called FortiGate 1 and the Branch FortiGate unit will be called FortiGate 2.

The steps include:

1. Creating redundant IPsec tunnels on FortiGate 1.
2. Configuring IP addresses and OSPF on FortiGate 1.
3. Configuring firewall addresses on FortiGate 1.
4. Configuring security policies on FortiGate 1.
5. Creating redundant IPsec tunnels for FortiGate 2.
6. Configuring IP addresses and OSPF on FortiGate 2.
7. Configuring firewall addresses on FortiGate 2.
8. Configuring security policies on FortiGate 2.

Creating redundant IPsec tunnels on FortiGate 1

1. Go to **VPN > IPsec Tunnels**.
2. Select **Create New**, name the primary tunnel and select **Custom VPN Tunnel (No Template)**.
3. Set the following:

Remote Gateway	Static IP Address
IP Address	FortiGate 2's wan1 IP
Local Interface	wan1 (the primary Internet-facing interface)
Pre-shared Key	Enter

4. Go to **VPN > IPsec Tunnels**.
5. Select **Create New**, name the secondary tunnel and select **Custom VPN Tunnel (No Template)**.
6. Set the following:

Remote Gateway	Static IP Address
IP Address	FortiGate 2's wan2 IP
Local Interface	wan2 (the secondary Internet-facing interface)
Pre-shared Key	Enter

Configuring IP addresses and OSPF on FortiGate 1

1. Go to **Network > Interfaces**.
2. Select the arrow for **wan1** to expand the list.

3. Edit the primary tunnel interface and create IP addresses.

IP	10.1.1.1
Remote IP	10.1.1.2

4. Select the arrow for **wan2** to expand the list.
5. Edit the secondary tunnel interface and create IP addresses.

IP	10.2.1.1
Remote IP	10.2.1.2

6. Go to **Network > OSPF** and enter the **Router ID** for FortiGate 1.
7. Select **Create New** in the **Area** section.
8. Add the backbone area of **0.0.0.0**.
9. Select **Create New** in the **Networks** section.
10. Create the networks and select **Area 0.0.0.0** for each one.
11. Select **Create New** in the **Interfaces** section.
12. Create primary and secondary tunnel interfaces.
13. Set a **Cost** of **10** for the primary interface and **100** for the secondary interface.

Configuring firewall addresses on FortiGate 1

1. Go to **Policy & Objects > Addresses**.
2. Create/Edit the subnets behind FortiGate 1 and FortiGate 2.
3. Create/Edit the primary and secondary interfaces of FortiGate 2.

Configuring security policies on FortiGate 1

1. Go to **Policy & Objects > IPv4 Policy**.
2. Create the four security policies required for both FortiGate 1's primary and secondary interfaces to connect to FortiGate 2's primary and secondary interfaces.

Creating redundant IPsec tunnels on FortiGate 2

1. Go to **VPN > IPsec Tunnels**.
2. Select **Create New**, name the primary tunnel and select **Custom VPN Tunnel (No Template)**.
3. Set the following:

Remote Gateway	Static IP Address
IP Address	FortiGate 1's wan1 IP
Local Interface	wan1 (the primary Internet-facing interface)
Pre-shared Key	Enter

4. Go to **VPN > IPsec Tunnels**.
5. Select **Create New**, name the secondary tunnel and select **Custom VPN Tunnel (No Template)**.
6. Set the following:

Remote Gateway	Static IP Address
IP Address	FortiGate 1's wan1 IP
Local Interface	wan2 (the secondary Internet-facing interface)
Pre-shared Key	Enter

Configuring IP addresses and OSPF on FortiGate 1

1. Go to **Network > Interfaces**.
2. Select the arrow for **wan1** to expand the list.
3. Edit the primary tunnel interface and create IP addresses.

IP	10.1.1.2
Remote IP	10.1.1.1

4. Select the arrow for **wan2** to expand the list.
5. Edit the secondary tunnel interface and create IP addresses.

IP	10.2.1.2
Remote IP	10.2.1.1

6. Go to **Network > OSPF** and enter the **Router ID** for FortiGate 2.
7. Select **Create New** in the **Area** section.
8. Add the backbone area of **0.0.0.0**.
9. Select **Create New** in the **Networks** section.
10. Create the networks and select **Area 0.0.0.0** for each one.
11. Select **Create New** in the **Interfaces** section.
12. Create primary and secondary tunnel interfaces.
13. Set a **Cost** of **10** for the primary interface and **100** for the secondary interface.

Configuring firewall addresses on FortiGate 2

1. Go to **Policy & Objects > Addresses**.
2. Create/Edit the subnets behind FortiGate 1 and FortiGate 2.
3. Create/Edit the primary and secondary interfaces of FortiGate 2.

Configuring security policies on FortiGate 2

1. Go to **Policy & Objects > IPv4 Policy**.
2. Create the four security policies required for both FortiGate 2's primary and secondary interfaces to connect to FortiGate 1's primary and secondary interfaces.

Results

1. Go to **Monitor > IPsec Monitor** to verify the statuses of both the primary and secondary IPsec VPN tunnels on FortiGate 1 and FortiGate 2.
2. Go to **Monitor > Routing Monitor**. Monitor to verify the routing table on FortiGate 1 and FortiGate 2. Type **OSPF** for the **Type** and select **Apply Filter** to verify the OSPF route.
3. Verify that traffic flows via the primary tunnel:
 - From a PC1 set to IP:10.20.1.100 behind FortiGate 1, run a tracer to a PC2 set to IP address 10.21.1.100 behind FortiGate 2 and vice versa.
 - From PC1, you should see that the traffic goes through 10.1.1.2 which is the primary tunnel interface IP set on FortiGate 2.
 - From PC2, you should see the traffic goes through 10.1.1.1 which is the primary tunnel interface IP set on FortiGate 1.
4. The VPN network between the two OSPF networks uses the primary VPN connection. Disconnect the wan1 interface and confirm that the secondary tunnel will be used automatically to maintain a secure connection.
5. Verify the IPsec VPN tunnel statuses on FortiGate 1 and FortiGate 2. Both FortiGates should show that primary tunnel is DOWN and secondary tunnel is UP.
6. Go to **Monitor > IPsec Monitor** to verify the status.
7. Verify the routing table on FortiGate 1 and FortiGate 2.
The secondary OSPF route (with cost = 100) appears on both FortiGate units.
8. Go to **Monitor > Routing Monitor**. Type **OSPF** for the **Type** and select **Apply Filter** to verify OSPF route.
9. Verify that traffic flows via the secondary tunnel:
 - From a PC1 set to IP:10.20.1.100 behind FortiGate 1, run a tracer to a PC2 set to IP:10.21.1.100 behind FortiGate 2 and vice versa.
 - From PC1, you should see that the traffic goes through 10.2.1.2 which is the secondary tunnel interface IP set on FortiGate 2.
 - From PC2, you should see the traffic goes through 10.2.1.1 which is the secondary tunnel interface IP set on FortiGate 1.

OSPF over dynamic IPsec

The following example shows how to create a dynamic IPsec VPN tunnel that allows OSPF.

Configuring IPsec on FortiGate 1

1. Go to **Dashboard** and enter the **CLI Console** widget
2. Create phase 1:

```
config vpn ipsec phase1-interface
  edit "dial-up"
    set type dynamic
    set interface "wan1"
    set mode-cfg enable
    set proposal 3des-shal
    set add-route disable
    set ipv4-start-ip 10.10.101.0
    set ipv4-end-ip 10.10.101.255
    set psksecret
  next
end
```

3. Create phase 2:

```
config vpn ipsec phase2-interface
  edit "dial-up-p2"
    set phase1name "dial-up"
    set proposal 3des-shal aes128-shal
  next
end
```

Configuring OSPF on FortiGate 1

1. Go to **Dashboard** and enter the **CLI Console** widget.
2. Create OSPF route.

```
config router ospf
  set router-id 172.20.120.22
  config area
    edit 0.0.0.0
    next
  end
  config network
    edit 1
      set prefix 10.10.101.0 255.255.255.0
    next
  end
  config redistribute "connected"
    set status enable
  end
  config redistribute "static"
    set status enable
  end
end
```

Adding policies on FortiGate 1

1. Go to **Policy & Objects > IPv4 Policy** and create a policy allowing OSPF traffic from **dial-up** to **port5**.
2. Go to **Policy & Objects > IPv4 Policy** and create a policy allowing OSPF traffic from **port5** to **dial-up** interfaces.

Configuring IPsec on FortiGate 2

1. Go to **Dashboard** and enter the **CLI Console** widget
2. Create phase 1:

```
config vpn ipsec phase1-interface
  edit "dial-up-client"
    set interface "wan1"
    set mode-cfg enable
    set proposal 3des-sha1
    set add-route disable
    set remote-gw 172.20.120.22
    set psksecret
  next
end
```

3. Create phase 2:

```
config vpn ipsec phase2-interface
  edit "dial-up-client"
    set phase1name "dial-up-client"
    set proposal 3des-sha1 aes128-sha1
    set auto-negotiate enable
  next
end
```

Configuring OSPF on FortiGate 2

1. Go to **Dashboard** and enter the **CLI Console** widget.
2. Create OSPF route.

```
config router ospf
  set router-id 172.20.120.15
  config area
    edit 0.0.0.0
    next
  end
  config network
    edit 1
      set prefix 10.10.101.0 255.255.255.0
    next
  end
  config redistribute "connected"
    set status enable
  end
  config redistribute "static"
    set status enable
  end
end
```

Adding policies on FortiGate 2

1. Go to **Policy & Objects > IPv4 Policy** and create a policy allowing OSPF traffic from **dial-up-client** to **port5**.
2. Go to **Policy & Objects > IPv4 Policy** and create a policy allowing OSPF traffic from **port5** to **dial-up-client** interfaces.

Verifying the tunnel is up

Go to **Monitor > IPsec Monitor** to verify that the tunnel is **Up**.

Results

1. From FortiGate 1, go to **Monitor > Routing Monitor** and verify that routes from FortiGate 2 were successfully advertised to FortiGate 1 via OSPF.
2. From FortiGate 1, go to **Dashboard**. Enter the CLI Console widget and type this command to verify OSPF neighbors:

```
get router info ospf neighbor
```

```
OSPF process 0:
Neighbor      ID Pri State Dead   Time      Address Interface
172.20.120.25 1  Full /      -    00:00:34 10.10.101.1 dial-up_0
```

3. From FortiGate 2, go to **Monitor > Routing Monitor** and verify that routes from FortiGate 1 were successfully advertised to FortiGate 2 via OSPF.
4. From FortiGate 2, go to **Dashboard**. Enter the CLI Console widget and type this command to verify OSPF neighbors:

```
get router info ospf neighbor
```

```
OSPF process 0:
Neighbor      ID Pri State Dead   Time      Address Interface
172.20.120.22 1  Full /      -    00:00:30 10.10.101.2 dial-up_client
```

BGP over dynamic IPsec

The following example shows how to create a dynamic IPsec VPN tunnel that allows BGP.

Configuring IPsec on FortiGate 1

1. Go to **Policy & Objects > Addresses** and select create new **Address**.

Name	Remote_loop_int
Type	Subnet
Subnet/IP Range	10.10.10.10
Interface	any

2. Create an **Address Group**.

Group Name	VPN_DST
Show in Address List	enable
Members	Remote_loop_int all

3. Go to **Dashboard** and enter the CLI Console widget.
4. Create phase 1:

```
config vpn ipsec phase1-interface
edit Dialup
    set type dynamic
    set interface wan1
    set mode aggressive
    set peertype one
    set mode-cfg enable
    set proposal 3des-sha1 aes128-sha1
    set peerid dial
    set assign-ip disable
    set psksecret
next
end
```

5. Create phase 2:

```
config vpn ipsec phase2-interface
edit dial_p2
    set phase1name Dialup
    set proposal 3des-sha1 aes128-sha1
    set src-addr-type name
    set dst-addr-type name
    set src-name all
    set dst-name VPN_DST
```



```

    next
end

```

Configuring BGP on FortiGate 1

1. Go to **Network > Interfaces** and create a Loopback interface.
2. Set **IP/Network Mask** to **20.20.20.20/255.255.255.255**.
3. Go to **Dashboard** and enter the CLI Console widget.
4. Create a BGP route.

```

config router bgp
    set as 100
    set router-id 1.1.1.1
    config neighbor
        edit 10.10.10.10
            set ebgp-enforce-multihop enable
            set remote-as 200
            set update-source loop
        next
    end
    config redistribute connected
        set status enable
    end
end

```

Adding policies on FortiGate 1

1. Go to **Policy & Objects > IPv4 Policy** and create a policy allowing BGP traffic from **Dialup** to **loop** interfaces.
2. Go to **Policy & Objects > IPv4 Policy** and create a policy allowing BGP traffic from **loop** to **Dialup** interfaces.

Configuring IPsec on FortiGate 2

1. Go to **Dashboard** and enter the CLI Console widget.
2. Create phase 1:

```

config vpn ipsec phase1-interface
    edit Dialup
        set interface wan1
        set mode aggressive
        set mode-cfg enable
        set proposal 3des-sha1 aes128-sha1
        set localid dial
        set remote-gw 172.20.120.22
        set assign-ip disable
        set psksecret
    next
end

```

3. Create phase 2:

```

config vpn ipsec phase2-interface
    edit dial_p2
        set phase1name Dialup
        set proposal 3des-sha1 aes128-sha1
        set keepalive enable
    next
end

```

Configuring BGP on FortiGate 2

1. Go to **Network > Interfaces** and create a Loopback interface.
2. Set **IP/Network Mask** to **10.10.10.10/255.255.255.255**.
3. Go to **Dashboard** and enter the **CLI Console** widget.
4. Create a BGP route.

```
config router bgp
  set as 200
  set router-id 1.1.1.2
  config neighbor
    edit 20.20.20.20
      set ebgp-enforce-multihop enable
      set remote-as 100
      set update-source loop
    next
  end
  config redistribute connected
    set status enable
  end
end
```

Adding policies on FortiGate 2

1. Go to **Policy & Objects > IPv4 Policy** and create a policy allowing BGP traffic from **Dialup** to **loop** interfaces.
2. Go to **Policy & Objects > IPv4 Policy** and create a policy allowing BGP traffic from **loop** to **Dialup** interfaces.

Adding a static route on FortiGate 2

Go to **Network > Static Routes** and add a route to the remote Loopback interface via Dialup interface.

Destination IP/Mask	20.20.20.20/255.255.255.255
Device	Dialup
Administrative Distance	10

Verifying the tunnel is up

Go to **Monitor > IPsec Monitor** to verify that the tunnel is **Up**.

Results

1. From FortiGate 1, go to **Monitor > Routing Monitor** and verify that routes from FortiGate 2 were successfully advertised to FortiGate 1 via BGP.
2. From FortiGate 1, go to **Dashboard**.
3. Enter the **CLI Console** widget and type this command to verify BGP neighbors:


```
get router info bgp summary
```
4. From FortiGate 2, go to **Monitor > Routing Monitor** and verify that routes from FortiGate 1 were successfully advertised to FortiGate 2 via BGP.

5. From FortiGate 2, go to **Dashboard**.
6. Enter the **CLI Console** widget and type this command to verify BGP neighbors:

```
get router info bgp summary
```

IPsec Auto-Discovery VPN (ADVPN)

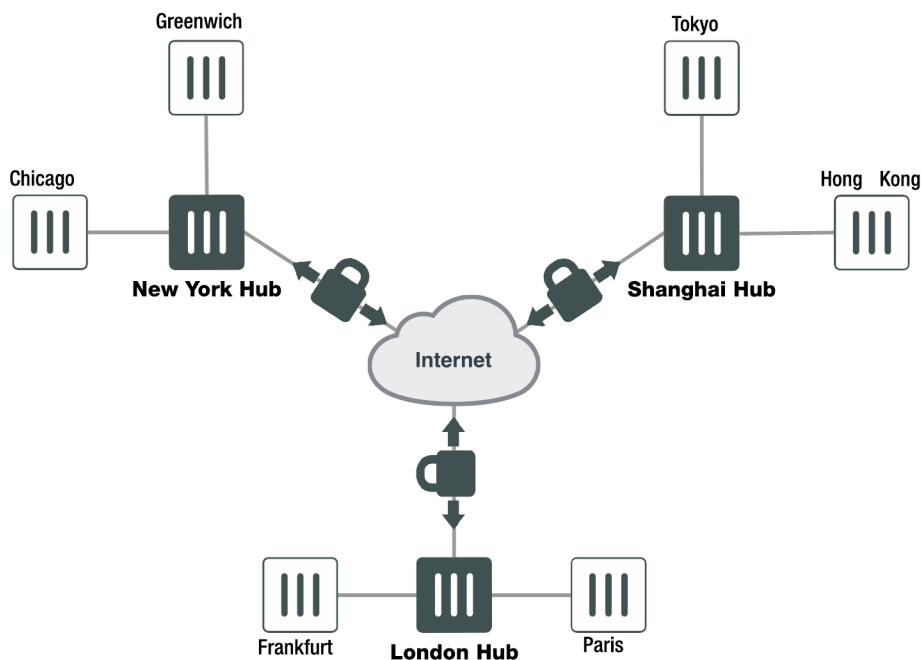
Consider a company that wants to provide direct secure (IPsec) connections between all of its offices in New York, Chicago, Greenwich, London, Paris, Frankfurt, Tokyo, Shanghai, and Hong Kong.

A straightforward solution is to create a full mesh of connections such that every site has eight IPsec configurations, one for each of the other sites. If there were ninety sites, that could still be done but now the configuration is becoming tedious, since every time a new site is added, N-1 other sites have to have their configuration updated.

An efficient and secure alternative is IPsec Auto-Discovery VPN (ADVPN), which allows a minimum amount of configuration per site but still allows direct IPsec connections to be made between every site. [RFC 7018](#) essentially describes this problem, along with some requirements for candidate solutions.

The ADVPN solution involves partitioning the sites into spokes and hubs such that a spoke has to have enough IPsec configuration to enable it to connect to at least one hub. A hub does not have specific configuration for each spoke, so the amount of configuration does not grow with the number of spokes that are connected to that hub. A hub to hub connection would typically involve both hubs having configuration for each other.

So, one possible partition for the original nine sites would be that Chicago and Greenwich would be spokes for the New York hub, Paris and Frankfurt would be spokes for the London hub, and Tokyo and Hong Kong would be spokes for the Shanghai hub:



Once a spoke has established a connection to its hub then initially IPsec traffic to another site transits via one or more hubs. For example, traffic from Chicago to Hong Kong would transit via the New York and Shanghai hubs. This transit traffic then triggers an attempt to create a more direct connection.

In FortiOS:

- Direct connections are only created between the two endpoints that want to exchange traffic (e.g. Chicago and Hong Kong); we do not create intermediate connections (say Chicago to Shanghai, or New York to Hong Kong) as a side-effect.
- Learning the peer subnets is done via a dynamic routing protocol running over the IPsec connections.
- Negotiation of the direct connections is done via IKE.
- Both PSK and certificate authentication is supported.

Example ADVPN configuration


Since dynamic routing with IPsec under FortiOS requires that an interface have an IP address, then for every site a unique IP address from some unused range is allocated. For example we'll assume that 10.100.0.0/16 is unused and so assign the IP addresses:

- | | | |
|------------------------|-----------------------|------------------------|
| • Chicago 10.100.0.4 | • London 10.100.0.2 | • Frankfurt 10.100.0.7 |
| • Greenwich 10.100.0.5 | • Shanghai 10.100.0.3 | • Hong Kong 10.100.0.8 |
| • New York 10.100.0.1 | • Paris 10.100.0.6 | • Tokyo 10.100.0.9 |

We'll assume that each site has one or more subnets that it protects that it wants to make available to the peers. For the purposes of exposition we'll assume there is only one subnet per site and they are allocated as:

- | | | |
|-------------------------|------------------------|-------------------------|
| • Chicago 10.0.4.0/16 | • London 10.0.2.0/24 | • Frankfurt 10.0.7.0/24 |
| • Greenwich 10.0.5.0/24 | • Shanghai 10.0.3.0/24 | • Hong Kong 10.0.8.0/24 |
| • New York 10.0.1.0/24 | • Paris 10.0.6.0/24 | • Tokyo 10.0.9.0/24 |

Our example network topology now looks like this:



The configuration in Chicago would be as follows:

```
config vpn ipsec phase1-interface
  edit "New York"
    set type static
    set interface wan1
    set remote-gw <New-York-IP-address>
    set psk <New-York-PSK>
    set auto-discovery-receiver enable
  next
end
```

The attribute `auto-discovery-receiver` indicates that this IPsec tunnel wishes to participate in an auto-discovery VPN. The IPsec interface would then have its IP assigned according to the Chicago address:

```
config system interface
  edit "New York"
    set ip 10.100.0.4/32
    set remote-ip 10.100.0.1
  next
end
```

RIP (for simplicity, you could use OSPF or BGP) is then configured to run on the IPsec interface and on the Chicago subnet (you could use `redistribute connected`, but we'll allow for the fact that there may be other subnets learned from another router on the 10.0.4.0/24 subnet):

```
config router rip
```

```
edit 1
    set prefix 10.100.0.0/16
next
edit 2
    set prefix 10.0.4.0/24
next
end
```

Other than the firewall policy and a minimal phase 2 configuration, this concludes the configuration for Chicago.

Each spoke would have a similar configuration.

The New York hub would have a dynamic phase 1 for its spoke connections, and two static phase 1s for its connections to the other hubs:

```
config vpn ipsec phase1-interface
edit "Spokes"
    set type dynamic
    set interface wan1
    set psk <New-York-PSK>
    set auto-discovery-sender enable
    set auto-discovery-psk enable
    set add-route disable
next
edit "London"
    set type static
    set interface wan1
    set psk <New-York-London-PSK>
    set auto-discovery-forwarder enable
next
edit "Shanghai"
    set type static
    set interface wan1
    set psk <New-York-Shanghai-PSK>
    set auto-discovery-forwarder enable
next
end
```

The 'Spokes' connection has `set auto-discovery-sender enable` to indicate that when IPsec traffic transits the hub it should optionally generate a message to the initiator of the traffic to indicate that it could perhaps establish a more direct connection. The `set add-route disable` ensures that IKE does not automatically add a route back over the spoke and instead leaves routing to a separately configured routing protocol.

The two inter-hub connections have `set auto-discovery-forwarder enable` to indicate that these connections can participate in the auto-discovery process. The interface IP addresses are assigned:

```
config system interface
edit "Spokes"
    set ip 10.100.0.1/32
    set remote-ip 10.100.0.254
next
edit "London"
    set ip 10.100.0.1/32
    set remote-ip 10.100.0.2
next
edit "Shanghai"
    set ip 10.100.0.1/32
```

```
        set remote-ip 10.100.0.3
    next
end
```

Following this, RIP is enabled on the relevant interfaces:

```
config router rip
    edit 1
        set prefix 10.100.0.0/16
    next
    edit 2
        set prefix 10.0.1.0/24
    next
end
```

A similar configuration would be used on the other two hubs.

Traffic flow and tunnel connection

With the configuration in place at all spokes and hubs, assuming all the spokes are connected to a hub, then Chicago would learn (via RIP) that the route to the Hong Kong subnet 10.0.8.0/24 is via its "New York" interface. If a device on the Chicago protected subnet (say 10.0.4.45) attempted to send traffic to the Hong Kong protected subnet (say 10.0.8.13) then it should flow over the New York interface to New York, which should then transmit it over the Shanghai tunnel to Shanghai, which should then send it over the dynamically negotiated Hong Kong tunnel to Hong Kong.

At the point when the traffic transits New York it should notice that the Chicago Spoke tunnel and the Shanghai tunnel have auto-discovery enabled, causing the New York hub to send a message via IKE to Chicago informing it that it may want to try and negotiate a direct connection for traffic from 10.0.4.45 to 10.0.8.13.

On receipt of this message, IKE on Chicago creates the (FortiOS-specific) IKE INFORMATIONAL SHORTCUT-QUERY message which contains the Chicago public IP address, the source IP of the traffic (10.0.4.45), the desired destination IP (10.0.8.13), and the PSK that should be used to secure any direct tunnel (if certificates are configured, it is assumed that they all share the same CA and so no additional authentication information is required). This message is sent via IKE to New York since routing indicates that New York is the best route to 10.0.8.13.

On receipt of the IKE INFORMATIONAL query, New York checks its routing table to see who owns 10.0.8.13. It finds that 10.0.8.13 should be routed via Shanghai, and since Shanghai is marked as an auto-discovery-forwarder then the query is forwarded.

Shanghai repeats the process, finds that 10.0.8.13 should be routed via its Hong Kong Spoke and so sends it to Hong Kong. Hong Kong checks 10.0.8.13, finds that it owns the subnet, so it remembers the Chicago public IP address (and PSK) and creates an IKE INFORMATIONAL reply message containing its external IP address. To work out where to send the IKE message, the FortiGate does a routing lookup for the original source IP (10.0.4.45), determines that the message should be routed via its Shanghai tunnel and so sends the reply back to Shanghai. The reply then makes its way back to Chicago following the reverse of the path that it used to arrive at Hong Kong.

When the reply makes it back to the Chicago initiator then it now knows the IP address of the Hong Kong device. Chicago now creates a new dynamic tunnel with the remote gateway as the Hong Kong public IP address and initiates an IKE negotiation (the dynamic tunnel name is auto-generated from the tunnel over which it performed the query; in this case it would be called 'New York_0').

This negotiation should succeed since Hong Kong is set up to expect an attempted negotiation from the Chicago public IP address. Once the negotiation succeeds, RIP will start to run on the newly created tunnels at Chicago and Hong Kong. This will update the routing on Chicago (and Hong Kong) so that the preferred route to 10.0.8.0 (10.0.4.0) is via the newly created tunnel rather than via the connection to New York (Shanghai).

Notes about ADVPN in FortiOS

- Auto-discovery is only supported by IKEv1.
- All Spokes must have an IP address that is routable from any other spoke; devices behind NAT are not currently supported.
- The feature requires the use of a dynamic routing protocol. There is no support for IKE handling routing.
- RIP is not a very scalable routing protocol. When there are more than a few spokes it would be advisable to use route summarization to avoid huge RIP updates. Better yet, use BGP instead of RIP.
- It is assumed that spokes will not be used to transit other spoke traffic, for example: traffic from Chicago to Tokyo would not transit an existing Chicago to Hong Kong tunnel even though that has a shorter hop count than a route via New York and Shanghai.
- There is no facility to allow you to filter which traffic that transits the hub should trigger the message sent to the initiator suggesting it create a direct connection. Currently any and all traffic will trigger it.

Logging and monitoring

This section provides some general logging and monitoring procedures for VPNs.

The following topics are included in this section:

Monitoring VPN connections

You can use the monitor to view activity on IPsec VPN tunnels and to start or stop those tunnels. The display provides a list of addresses, proxy IDs, and timeout information for all active tunnels.

Monitoring connections to remote peers

The list of tunnels provides information about VPN connections to remote peers that have static IP addresses or domain names. You can use this list to view status and IP addressing information for each tunnel configuration. You can also start and stop individual tunnels from the list.

To view the list of static-IP and dynamic-DNS tunnels go to **Monitor > IPsec Monitor**.

Monitoring dialup IPsec connections

The list of dialup tunnels provides information about the status of tunnels that have been established for dialup clients. The list displays the IP addresses of dialup clients and the names of all active tunnels. The number of tunnels shown in the list can change as dialup clients connect and disconnect.

To view the list of dialup tunnels go to **Monitor > IPsec Monitor**.

If you take down an active tunnel while a dialup client such as FortiClient is still connected, FortiClient will continue to show the tunnel connected and idle. The dialup client must disconnect before another tunnel can be initiated.

The list of dialup tunnels displays the following statistics:

- The Name column displays the name of the tunnel.
- The meaning of the value in the Remote gateway column changes, depending on the configuration of the network at the far end:
 - When a FortiClient dialup client establishes a tunnel, the Remote gateway column displays either the public IP address and UDP port of the remote host device (on which the FortiClient Endpoint Security application is installed), or if a NAT device exists in front of the remote host, the Remote gateway column displays the public IP address and UDP port of the remote host.
 - When a FortiGate dialup client establishes a tunnel, the Remote gateway column displays the public IP address and UDP port of the FortiGate dialup client.
- The Username column displays the peer ID, certificate name, or XAuth user name of the dialup client (if a peer ID, certificate name, or XAuth user name was assigned to the dialup client for authentication purposes).
- The Timeout column displays the time before the next key exchange. The time is calculated by subtracting the time elapsed since the last key exchange from the keylife.
- The Proxy ID Source column displays the IP addresses of the hosts, servers, or private networks behind the FortiGate unit. A network range may be displayed if the source address in the security encryption policy was expressed as a range of IP addresses.

- The meaning of the value in the Proxy ID Destination column changes, depending on the configuration of the network at the far end:
 - When a FortiClient dialup client establishes a tunnel:
 - If VIP addresses are not used and the remote host connects to the Internet directly, the Proxy ID Destination field displays the public IP address of the Network Interface Card (NIC) in the remote host.
 - If VIP addresses are not used and the remote host is behind a NAT device, the Proxy ID Destination field displays the private IP address of the NIC in the remote host.
 - If VIP addresses were configured (manually or through FortiGate DHCP relay), the Proxy ID Destination field displays either the VIP address belonging to a FortiClient dialup client, or a subnet address from which VIP addresses were assigned.
- When a FortiGate dialup client establishes a tunnel, the Proxy ID Destination field displays the IP address of the remote private network.

VPN event logs

You can configure the FortiGate unit to log VPN events. For IPsec VPNs, Phase 1 and Phase 2 authentication and encryption events are logged. For information about how to interpret log messages, see the [FortiGate Log Message Reference](#).

Logging VPN events

1. Go to **Log & Report > Log Settings**.
2. Verify that the **VPN activity event** option is selected.
3. Select **Apply**.

Viewing event logs

1. Go to **Log & Report > VPN Events**.
2. Select the **Log location**.

Sending tunnel statistics to FortiAnalyzer

By default, logged events include tunnel-up and tunnel-down status events. Other events, by default, will appear in the FortiAnalyzer report as "No Data Available". More accurate results require logs with `action=tunnel-stats`, which is used in generating reports on the FortiAnalyzer (rather than the tunnel-up and tunnel-down event logs). The FortiGate does not, by default, send `tunnel-stats` information.

To allow VPN `tunnel-stats` to be sent to FortiAnalyzer, configure the FortiGate unit as follows using the CLI:

```
config system settings
    set vpn-stats-log ipsec ssl
    set vpn-stats-period 300
end
```

Troubleshooting

This section contains tips to help you with some common challenges of IPsec VPNs.

A VPN connection has multiple stages that can be confirmed to ensure the connection is working properly. It is easiest to see if the final stage is successful first since if it is successful the other stages will be working properly. Otherwise, you will need to work back through the stages to see where the problem is located.

When a VPN connection is properly established, traffic will flow from one end to the other as if both ends were physically in the same place. If you can determine the connection is working properly then any problems are likely problems with your applications.

On some FortiGate units, such as the FortiGate 94D, you cannot ping over the IPsec tunnel without first setting a source-IP. In this scenario, you must assign an IP address to the virtual IPsec VPN interface. Anything sourced from the FortiGate going over the VPN will use this IP address.

If the egress/outgoing interface (determined by kernel route) has an IP address, then use the IP address of the egress/outgoing interface. Otherwise, use the IP address of the first interface from the interface list (that has an IP address).

The first diagnostic command worth running, in any IPsec VPN troubleshooting situation, is the following:

```
diagnose vpn tunnel list
```

This command is very useful for gathering statistical data such as the number of packets encrypted versus decrypted, the number of bytes sent versus received, the SPI identifier, etc. This kind of information in the resulting output can make all the difference in determining the issue with the VPN.

Another appropriate diagnostic command worth trying is:

```
diagnose debug flow
```

This command will inform you of any lack of firewall policy, lack of forwarding route, and of policy ordering issues.

Common IPsec VPN problems

The most common IPsec VPN issues are listed below. Please read thoroughly and note that, although the list is extensive, it is not exhaustive.

This section includes support for the following:

- [Failed VPN connection attempts](#)
- [Debug output table](#)
- [The options to configure policy-based IPsec VPN are unavailable](#)
- [The VPN tunnel goes down frequently](#)
- [The pre-shared key does not match \(PSK mismatch error\)](#)
- [The SA proposals do not match \(SA proposal mismatch\)](#)
- [Pre-existing IPsec VPN tunnels need to be cleared](#)
- [Other potential VPN issues](#)

Failed VPN connection attempts

If your VPN fails to connect, check the following:

- Ensure that the **pre-shared keys** match exactly (see [The pre-shared key does not match \(PSK mismatch error\)](#) below).
- Ensure that both ends use the same P1 and P2 proposal settings (see [The SA proposals do not match \(SA proposal mismatch\)](#) below).
- Ensure that you have allowed inbound and outbound traffic for all necessary network services, especially if services such as DNS or DHCP are having problems.
- Check that a static route has been configured properly to allow routing of VPN traffic.

If you are still unable to connect to the VPN tunnel, run the following diagnostic command in the CLI:

```
diagnose debug application ike -1
diagnose debug enable
```

The resulting output may indicate where the problem is occurring. When you are finished, disable the diagnostics by using the following command:

```
diagnose debug reset
diagnose debug disable
```

View the table below for some assistance in analyzing the debug output.

Debug output table

Problem	Debug output	Common causes	Common solutions
Tunnel is not coming up	Error: negotiation failure	IPsec configuration mismatch	Check phase 1 and 2 settings
	Error: no SA proposal chosen	IPsec configuration mismatch	Check phase 1 and 2 settings
	FortiGate using the wrong VPN	Missing or wrong local ID	If there are more than one pre-shared key dial-up VPN with the same local gateway, use aggressive mode and different local IDs
	Error: connection expiring due to XAUTH failure	Wrong username, password, or user group	Check user credentials and user group configuration
	Error: peer has not completed XAUTH exchange	XAuth is disabled in the client	Fix the client's XAuth configuration
Tunnel is bouncing	DPD packets lost	ISP issue	Check the ISP connection

Problem	Debug output	Common causes	Common solutions
Tunnel is up but traffic does not go through	Error: No matching IPsec selector, drop	Quick mode selector mismatch	Fix the quick mode selector
		NAT is enabled	Disable NAT in the firewall policy
	Traffic is not routed to the tunnel	Route or firewall policy misconfiguration	Route-based: traffic must be routed to IPsec virtual interface Policy-based: traffic must match a firewall policy with action set to IPSEC

The options to configure policy-based IPsec VPN are unavailable

Go to **System > Feature Visibility**. Select **Show More** and turn on **Policy-based IPsec VPN**.

The VPN tunnel goes down frequently

If your VPN tunnel goes down often, check the Phase 2 settings and either increase the **Keylife** value or enable **Autokey Keep Alive**.

The pre-shared key does not match (PSK mismatch error)

It is possible to identify a PSK mismatch using the following combination of CLI commands:

```
diag vpn ike log filter name <phase1-name>
diag debug app ike -1
diag debug enable
```

This will provide you with clues as to any PSK or other proposal issues. If it is a PSK mismatch, you should see something similar to the following output:

```
ike 0:TRX:322: PSK auth failed: probable pre-shared key mismatch
ike Negotiate SA Error:
```

The SA proposals do not match (SA proposal mismatch)

The most common problem with IPsec VPN tunnels is a mismatch between the proposals offered between each party. Without a match and proposal agreement, Phase 1 can never establish. Use the following command to show the proposals presented by both parties.

```
diag debug app ike -1
diag debug enable
```

The resulting output should include something similar to the following, where **blue** represents the remote VPN device, and **green** represents the local FortiGate.

```
responder received SA_INIT msg
incoming proposal:
proposal id = 1:
  protocol = IKEv2:
    encapsulation = IKEv2/none
    type=ENCR, val=AES_CBC (key_len = 256)
    type=INTEGR, val=AUTH_HMAC_SHA_96
    type=PRF, val=PRF_HMAC_SHA
    type=DH_GROUP, val=1536.
```

```
proposal id = 2:
  protocol = IKEv2:
    encapsulation = IKEv2/none
    type=ENCR, val=3DES_CBC
    type=INTEGR, val=AUTH_HMAC_SHA_2_256_128
    type=PRF, val=PRF_HMAC_SHA2_256
    type=DH_GROUP, val=1536.
proposal id = 1:
  protocol = IKEv2:
    encapsulation = IKEv2/none
    type=ENCR, val=AES_CBC (key_len = 128)
    type=INTEGR, val=AUTH_HMAC_SHA_96
    type=PRF, val=PRF_HMAC_SHA
    type=DH_GROUP, val=1536.
```

Pre-existing IPsec VPN tunnels need to be cleared

Should you need to clear an IKE gateway, use the following commands:

```
diagnose vpn ike restart
diagnose vpn ike gateway clear
```

Other potential VPN issues

- Ensure that your FortiGate unit is in NAT/Route mode, rather than Transparent.
- Check your NAT settings, enabling NAT traversal in the Phase 1 configuration while disabling NAT in the security policy. You might need to pin the PAT/NAT session table, or use some of kind of NAT-T keepalive to avoid the expiration of your PAT/NAT translation.
- Ensure that both ends of the VPN tunnel are using Main mode, unless multiple dial-up tunnels are being used.
- Remove any Phase 1 or Phase 2 configurations that are not in use. If a duplicate instance of the VPN tunnel appears on the IPsec Monitor, reboot your FortiGate unit to try and clear the entry.
- If you have multiple dial-up IPsec VPNs, ensure that the peer ID is configured properly on the FortiGate and that clients have specified the correct local ID. Furthermore, in circumstances where multiple remote dialup VPN tunnels exist, each tunnel must have a peer ID set.
- If you are using FortiClient, ensure that your version is compatible with the FortiGate firmware by reading the FortiOS Release Notes.
- If you are using Perfect Forward Secrecy (PFS), ensure that it is used on both peers. You can use the `diagnose vpn tunnel list` command to troubleshoot this.
- Ensure that the **Quick Mode selectors** are correctly configured. If part of the setup currently uses firewall addresses or address groups, try changing it to either specify the IP addresses or use an expanded address range. This is especially useful if the remote endpoint is not a FortiGate device.
- If XAUTH is enabled, ensure that the settings are the same for both ends, and that the FortiGate unit is set to **Enable as Server**.
- Check IPsec VPN Maximum Transmission Unit (MTU) size. A 1500 byte MTU is going to exceed the overhead of the ESP-header, including the additional ip_header, etc. You can use the `diagnose vpn tunnel list` command to troubleshoot this.
- If your FortiGate unit is behind a NAT device, such as a router, configure port forwarding for UDP ports 500 and 4500.

Troubleshooting connection issues

The following section includes troubleshooting suggestions related to:

- [LAN interface connection](#)
- [Dialup connection](#)
- [Troubleshooting VPN connections](#)
- [Troubleshooting invalid ESP packets using Wireshark](#)
- [Attempting hardware offloading beyond SHA1](#)
- [Check Phase 1 proposal settings](#)
- [Check your routing](#)
- [Try enabling XAuth](#)

LAN interface connection

To confirm whether a VPN connection over LAN interfaces has been configured correctly, issue a ping or traceroute command on the network behind the FortiGate unit to test the connection to a computer on the remote network. If the connection is properly configured, a VPN tunnel will be established automatically when the first data packet destined for the remote network is intercepted by the FortiGate unit.

If the ping or traceroute fail, it indicates a connection problem between the two ends of the tunnel. This may or may not indicate problems with the VPN tunnel. You can confirm this by going to **Monitor > IPsec Monitor** where you will be able to see your connection. A green arrow means the tunnel is up and currently processing traffic. A red arrow means the tunnel is not processing traffic, and this VPN connection has a problem.

If the connection has problems, see [Troubleshooting VPN connections on page 1833](#).

Dialup connection

A dialup VPN connection has additional steps. To confirm that a VPN between a local network and a dialup client has been configured correctly, at the dialup client, issue a ping command to test the connection to the local network. The VPN tunnel initializes when the dialup client attempts to connect.

If the ping or traceroute fail, it indicates a connection problem between the two ends of the tunnel. This may or may not indicate problems with the VPN tunnel, or dialup client. As with the LAN connection, confirm the VPN tunnel is established by checking **Monitor > IPsec Monitor**.

Troubleshooting VPN connections

If you have determined that your VPN connection is not working properly through [Troubleshooting on page 1829](#), the next step is to verify that you have a phase2 connection.

If traffic is not passing through the FortiGate unit as you expect, ensure the traffic does not contain IPcomp packets (IP protocol 108, RFC 3173). FortiGate units do not allow IPcomp packets, they compress packet payload, preventing it from being scanned.

Testing Phase 1 and 2 connections is a bit more difficult than testing the working VPN. This is because they require diagnose CLI commands. These commands are typically used by Fortinet customer support to discover more information about your FortiGate unit and its current configuration.

Before you begin troubleshooting, you must:

- Configure FortiGate units on both ends for interface VPN
- Record the information in your VPN Phase 1 and Phase 2 configurations - for our example here the remote IP address is 10.11.101.10 and the names of the phases are Phase 1 and Phase 2
- Install a telnet or SSH client such as putty that allows logging of output
- Ensure that the admin interface supports your chosen connection protocol so you can connect to your FortiGate unit admin interface.

For this example, default values were used unless stated otherwise.

Obtaining diagnose information for the VPN connection - CLI

1. Log into the CLI as admin with the output being logged to a file.
2. Stop any diagnose debug sessions that are currently running with the CLI command
`diagnose debug disable`
3. Clear any existing log-filters by running
`diagnose vpn ike log-filter clear`
4. Set the log-filter to the IP address of the remote computer (10.11.101.10). This filters out all VPN connections except ones to the IP address we are concerned with. The command is
`diagnose vpn ike log-filter dst-addr4 10.11.101.10.`
5. Set up the commands to output the VPN handshaking. The commands are:
`diagnose debug app ike 255`
`diagnose debug enable`

6. Have the remote FortiGate initiate the VPN connection in the web-based manager by going to **VPN > IPsec Tunnels** and selecting **Bring up**.

This makes the remote FortiGate the initiator and the local FortiGate becomes the responder. Establishing the connection in this manner means the local FortiGate will have its configuration information as well as the information the remote computer sends. Having both sets of information locally makes it easier to troubleshoot your VPN connection.

7. Watch the screen for output, and after roughly 15 seconds enter the following CLI command to stop the output.
`diagnose debug disable`
8. If needed, save the log file of this output to a file on your local computer. Saving the output to a file can make it easier to search for a particular phrase, and is useful for comparisons.

Troubleshooting a Phase 1 VPN connection

Using the output from [Obtaining diagnose information for the VPN connection - CLI](#), search for the word `proposal` in the output. It may occur once indicating a successful connection, or it will occur two or more times for an unsuccessful connection — there will be one proposal listed for each end of the tunnel and each possible combination in their settings. For example if 10.11.101.10 selected both Diffie-Hellman Groups 1 and 5, that would be at least 2 proposals set.

A successful negotiation proposal will look similar to

```
IPsec SA connect 26 10.12.101.10->10.11.101.10:500
config found
created connection: 0x2f55860 26 10.12.101.10->10.11.101.10:500
```

```

IPsec SA connect 26 10.12.101.10->10.11.101.10:500 negotiating
no suitable ISAKMP SA, queuing quick-mode request and initiating ISAKMP SA negotiation
initiator: main mode is sending 1st message...
cookie 3db6afe559e3df0f/0000000000000000
out [encryption]
sent IKE msg (ident=ilsend): 10.12.101.10:500->10.11.101.10:500, len=264,
    id=3db6afe559e3df0f/0000000000000000
diaike 0: comes 10.12.101.1:500->10.11.101.1:500,ifindex=26...

```

Note the phrase “initiator: main mode is sending 1st message...” which shows you the handshake between the ends of the tunnel is in progress. Initiator shows the remote unit is sending the first message.

Troubleshooting invalid ESP packets using Wireshark

The following section provides information to help debug an encryption key mismatch. The ESP packet invalid error is due to an encryption key mismatch after a VPN tunnel has been established. When an IPsec VPN tunnel is up, but traffic is not able to pass through the tunnel, Wireshark (or an equivalent program) can be used to determine whether there is an encryption mismatch. A mismatch could occur for many reasons, one of the most common is the instability of an ISP link (ADSL, Cable), or it could effectively be any device in the physical connection.

The following information is required to troubleshoot the problem.

- Take a packet sniffer trace on both FortiGates.
- Run the `diag vpn tunnel list` command a few times on both FortiGates when generating traffic that will pass through the tunnel.

In the following example, the error message was seen on the recipient FortiGate:

```

date=2010-12-28 time=18:19:35 devname=Kosad_VPN device_id=FG300B3910600118 log_
id=0101037132 type=event subtype=ipsec pri=critical vd="root" msg="IPsec ESP" action="error" rem_
ip=180.87.33.2 loc_ip=121.133.8.18 rem_port=32528 loc_port=4500 out_intf="port2"
cookies="88d40f65d555ccaf/05464e20e4afc835" user="N/A" group="N/A" xauth_user="N/A" xauth_
group="N/A" vpn_tunnel="fortinet_0" status=esp_error error_num=Invalid ESP packet detected (HMAC
validation failed). spi=c32b09f7 seq=00000012

```

This is the output of the command `diag vpn tunnel list` on the FortiGate:

```

inet ver=1 serial=2 192.168.1.205:4500->121.133.8.18:4500 lgwy=dyn tun=intf mode=auto bound_if=4
proxyid_num=1 child_num=0 refcnt=7 ilast=0 olast=0
stat: rxp=41 txp=56 rxb=4920 txb=3360
dpd: mode=active on=1 idle=5000ms retry=3 count=0 seqno=696
natt: mode=keepalive draft=32 interval=10 remote_port=4500
proxyid=P2_60C_Fortinet proto=0 sa=1 ref=2 auto_negotiate=0 serial=1 src:
0:182.40.101.0/255.255.255.0:0
dst: 0:100.100.100.0/255.255.255.0:0
SA: ref=3 options=0000000d type=00 soft=0 mtu=1428 expire=1106 replaywin=0 seqno=15
life: type=01 bytes=0/0 timeout=1777/1800
dec: spi=29a26eb6 esp=3des key=24 bf25e69df90257f64c55dda4069f01834cd0382fe4866ff2
ah=sha1 key=20 38b2600170585d2dfa646caed5bc86d920aed7ff
enc: spi=c32b09f7 esp=3des key=24 0abd3c70032123c3369a6f225a385d30f0b2fb1cd9687ec8
ah=sha1 key=20 214d8e717306dffceec3760464b6e8edb436c6

```

This is the packet capture from the FortiGate:

No.	Time	Source	Destination	Protocol	Info
39	48.222853	192.168.1.205	121.133.8.18	ISAKMP	Informational
40	48.348006	121.133.8.18	192.168.1.205	ISAKMP	Informational
41	48.348899	192.168.1.205	121.133.8.18	ISAKMP	Informational
42	48.738105	121.133.8.18	192.168.1.205	ISAKMP	Informational
43	51.607241	192.168.1.205	121.133.8.18	ESP	ESP (SPI=0xc22b09f7)
44	51.731181	192.168.1.205	121.133.8.18	ISAKMP	Informational
45	51.731388	121.133.8.18	192.168.1.205	ISAKMP	Informational
46	51.731226	192.168.1.205	121.133.8.18	ISAKMP	Informational
47	54.149077	121.133.8.18	192.168.1.205	ISAKMP	Informational
48	57.187934	192.168.1.205	121.133.8.18	ESP	ESP (SPI=0xc22b09f7)
49	59.142953	192.168.1.205	121.133.8.18	ISAKMP	Informational
50	59.138415	121.133.8.18	192.168.1.205	ISAKMP	Informational
51	59.159274	192.168.1.205	121.133.8.18	ISAKMP	Informational
52	59.544269	121.133.8.18	192.168.1.205	ISAKMP	Informational
53	62.688010	192.168.1.205	121.133.8.18	ESP	ESP (SPI=0xc22b09f7)
54	64.542991	192.168.1.205	121.133.8.18	ISAKMP	Informational
55	64.797978	121.133.8.18	192.168.1.205	ISAKMP	Informational
56	64.798749	192.168.1.205	121.133.8.18	ISAKMP	Informational
57	64.957383	121.133.8.18	192.168.1.205	ISAKMP	Informational

How to verify if the original packet has been encrypted correctly

To verify, it is necessary to decrypt the ESP packet using Wireshark. Open the packet capture that is taken from initiator FortiGate using Wireshark. Go to **Edit > Preferences**, expand **Protocol** and look for **ESP**. Select **"Attempt to detect/decode encrypted ESP payloads"**, and fill in the information for the encryption algorithm and the keys. This information can be obtained from the output of the command `diag vpn tunnel list`.

If the packet was encrypted correctly using the correct key, then the decryption will be successful and it will be possible to see the original package as shown below:

13	16.199042	121.133.8.18	192.168.1.205	ISAKMP	Informational
14	16.836503	192.168.1.205	100.100.100.202	ICMP	Echo (ping) request
15	21.192060	192.168.1.205	121.133.8.18	ISAKMP	Informational
16	21.223916	121.133.8.18	192.168.1.205	ISAKMP	Informational
17	21.224790	192.168.1.205	121.133.8.18	ISAKMP	Informational
18	21.392073	121.133.8.18	192.168.1.205	ISAKMP	Informational
19	24.187008	192.168.1.205	100.100.100.202	ICMP	Echo (ping) request
20	26.392870	192.168.1.205	121.133.8.18	ISAKMP	Informational
21	26.746759	121.133.8.18	192.168.1.205	ISAKMP	Informational
22	26.747595	192.168.1.205	121.133.8.18	ISAKMP	Informational
23	26.999062	121.133.8.18	192.168.1.205	ISAKMP	Informational
24	29.687170	192.168.1.205	100.100.100.202	ICMP	Echo (ping) request
25	31.092680	192.168.1.205	121.133.8.18	ISAKMP	Informational
26	32.143504	121.133.8.18	192.168.1.205	ISAKMP	Informational
27	32.143590	192.168.1.205	121.133.8.18	ISAKMP	Informational
28	32.288780	121.133.8.18	192.168.1.205	ISAKMP	Informational
29	35.182775	192.168.1.205	100.100.100.202	ICMP	Echo (ping) request
30	37.282867	192.168.1.205	121.133.8.18	ISAKMP	Informational
31	37.336742	121.133.8.18	192.168.1.205	ISAKMP	Informational

Repeat the decryption process for the packet capture from the recipient firewall. If the decryption failed using the same key, the packet may be corrupted and the interface should then be checked for CRC or packet errors.

Attempting hardware offloading beyond SHA1

If you are trying to off-load VPN processing to a network processing unit (NPU), remember that only SHA1 authentication is supported. For high levels of authentication such as SHA256, SHA384, and SHA512 hardware offloading is not an option—all VPN processing must be done in software—unless using an NP6 (although the NP4lite variation also supports SHA256, SHA384, and SHA512).

Enable/disable IPsec ASIC-offloading

Much like NPU-offload in IKE phase1 configuration, you can enable or disable the usage of ASIC hardware for IPsec Diffie-Hellman key exchange and IPsec ESP traffic. By default hardware offloading is used. For debugging purposes, sometimes it is best for all the traffic to be processed by software.

```
config sys global
    set ipsec-asic-offload [enable | disable]
end
```

Check Phase 1 proposal settings

Ensure that both sides have at least one Phase 1 proposal in common. Otherwise they will not connect. If there are many proposals in the list, this will slow down the negotiating of Phase 1. If its too slow, the connection may timeout before completing. If this happens, try removing some of the unused proposals.

NPU offloading is supported when the local gateway is a loopback interface.

Check your routing

If routing is not properly configured with an entry for the remote end of the VPN tunnel, traffic will not flow properly. You may need static routes on both ends of the tunnel. If routing is the problem, the proposal will likely setup properly but no traffic will flow.

Try enabling XAuth

If one end of an attempted VPN tunnel is using XAuth and the other end is not, the connection attempt will fail. The log messages for the attempted connection will not mention XAuth is the reason, but when connections are failing it is a good idea to ensure both ends have the same XAuth settings. If you do not know the other end's settings enable or disable XAuth on your end to see if that is the problem.

General troubleshooting tips

Most connection failures are due to a configuration mismatch between the FortiGate unit and the remote peer. In general, begin troubleshooting an IPsec VPN connection failure as follows:

1. Ping the remote network or client to verify whether the connection is up. See [General troubleshooting tips on page 1837](#).
2. Traceroute the remote network or client. If DNS is working, you can use domain names. Otherwise use IP addresses.
3. Check the routing behind the dialup client. Routing problems may be affecting DHCP. If this appears to be the case, configure a DHCP relay service to enable DHCP requests to be relayed to a DHCP server on or behind the FortiGate server.
4. Verify the configuration of the FortiGate unit and the remote peer. Check the following IPsec parameters:
 - The mode setting for ID protection (main or aggressive) on both VPN peers must be identical.
 - The authentication method (preshared keys or certificates) used by the client must be supported on the FortiGate unit and configured properly.
 - If preshared keys are being used for authentication purposes, both VPN peers must have identical preshared keys.
 - The remote client must have at least one set of Phase 1 encryption, authentication, and Diffie-Hellman settings that match corresponding settings on the FortiGate unit.
 - Both VPN peers must have the same NAT traversal setting (enabled or disabled).
 - The remote client must have at least one set of Phase 2 encryption and authentication algorithm settings that match the corresponding settings on the FortiGate unit.
 - If you are using manual keys to establish a tunnel, the **Remote SPI** setting on the FortiGate unit must be identical to the **Local SPI** setting on the remote peer, and vice versa.
5. To correct the problem, see the following table.

VPN troubleshooting tips

Configuration problem	Correction
Mode settings do not match.	Select complementary mode settings. See Phase 1 parameters on page 1655 .

Configuration problem	Correction
Peer ID or certificate name of the remote peer or dialup client is not recognized by FortiGate VPN server.	<p>Check Phase 1 configuration. Depending on the Remote Gateway and Authentication Method settings, you have a choice of options to authenticate FortiGate dialup clients or VPN peers by ID or certificate name (see Phase 1 parameters on page 1655).</p> <p>If you are configuring authentication parameters for FortiClient dialup clients, refer to the Authenticating FortiClient Dialup Clients Technical Note.</p>
Preshared keys do not match.	Reenter the preshared key. See Phase 1 parameters on page 1655 .
Phase 1 or Phase 2 key exchange proposals are mismatched.	Make sure that both VPN peers have at least one set of proposals in common for each phase. See Phase 1 parameters on page 1655 and Phase 2 parameters on page 1675 .
NAT traversal settings are mismatched.	Select or clear both options as required. See Phase 1 parameters on page 1655 and Phase 1 parameters on page 1655 .

A word about NAT devices

When a device with NAT capabilities is located between two VPN peers or a VPN peer and a dialup client, that device must be NAT traversal (NAT-T) compatible for encrypted traffic to pass through the NAT device. For more information, see [Phase 1 parameters on page 1655](#).

Troubleshooting L2TP and IPsec

This section describes some checks and tools you can use to resolve issues with L2TP-over-IPsec VPNs.

This section includes:

- [Quick checks](#)
- [Mac OS X and L2TP](#)
- [Setting up logging](#)
- [Using the FortiGate unit debug commands](#)

Quick checks

The table below is a list of common L2TP over IPsec VPN problems and the possible solutions.

Problem	What to check
IPsec tunnel does not come up.	<p>Check the logs to determine whether the failure is in Phase 1 or Phase 2.</p> <p>Check the settings, including encapsulation setting, which must be transport-mode.</p> <p>Check the user password.</p> <p>Confirm that the user is a member of the user group assigned to L2TP.</p> <p>On the Windows PC, check that the IPsec service is running and has not been disabled. See Troubleshooting L2TP and IPsec on page 1838.</p>
Tunnel connects, but there is no communication.	<p>Did you create an ACCEPT security policy from the public network to the protected network for the L2TP clients? See Troubleshooting L2TP and IPsec on page 1838.</p>

Mac OS X and L2TP

FortiOS allows L2TP connections with empty AVP host names and therefore Mac OS X L2TP connections can connect to the FortiGate.

Prior to FortiOS 4.0 MR3, FortiOS refused L2TP connections with empty AVP host names in compliance with [RFC 2661](#) and [RFC 3931](#).

Setting up logging

L2TP logging must be enabled to record L2TP events. Alert email can be configured to report L2TP errors.

Configuring FortiGate logging for L2TP over IPsec

1. Go to **Log & Report > Log Settings**.
2. Select **Event Log**.
3. Select the **VPN activity event** check box.
4. Select **Apply**.

Viewing FortiGate logs

1. Go to **Log & Report > VPN Events**.
2. Select the **Log location** if required.
3. After each attempt to start the L2TP over IPsec VPN, select **Refresh** to view logged events.

Using the FortiGate unit debug commands

Viewing debug output for IKE and L2TP

1. Start an SSH or Telnet session to your FortiGate unit.
2. Enter the following CLI commands

```
diagnose debug application ike -1
diagnose debug application l2tp -1
diagnose debug enable
```

3. Attempt to use the VPN and note the debug output in the SSH or Telnet session.
4. Enter the following command to reset debug settings to default:

```
diagnose debug reset
```

Using the packet sniffer

1. Start an SSH or Telnet session to your FortiGate unit.
2. Enter the following CLI command

```
diagnose sniffer packet any icmp 4
```

3. Attempt to use the VPN and note the debug output.
4. Enter Ctrl-C to end sniffer operation.

Typical L2TP over IPsec session startup log entries - raw format

```
2010-01-11 16:39:58 log_id=0101037127 type=event subtype=ipsec pri=notice vd="root" msg="progress IPsec Phase 1" action="negotiate" rem_ip=172.20.120.151 loc_ip=172.20.120.141 rem_port=500 loc_port=500 out_intf="port1" cookies="5f6dalc0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A" xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1" status=success init=remote mode=main dir=outbound stage=1 role=responder result=OK
```

```
2010-01-11 16:39:58 log_id=0101037127 type=event subtype=ipsec pri=notice vd="root" msg="progress IPsec Phase 1" action="negotiate" rem_ip=172.20.120.151 loc_ip=172.20.120.141 rem_port=500 loc_port=500 out_intf="port1" cookies="5f6dalc0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A" xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1" status=success init=remote mode=main dir=outbound stage=2 role=responder result=OK
```

```
2010-01-11 16:39:58 log_id=0101037127 type=event subtype=ipsec pri=notice vd="root" msg="progress IPsec Phase 1" action="negotiate" rem_ip=172.20.120.151 loc_ip=172.20.120.141 rem_port=500 loc_port=500 out_intf="port1" cookies="5f6dalc0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A" xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1" status=success init=remote mode=main dir=inbound stage=3 role=responder result=DONE
```

```
2010-01-11 16:39:58 log_id=0101037127 type=event subtype=ipsec pri=notice vd="root" msg="progress IPsec Phase 1" action="negotiate" rem_ip=172.20.120.151 loc_ip=172.20.120.141 rem_port=500 loc_port=500 out_intf="port1" cookies="5f6dalc0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A" xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1_0" status=success init=remote mode=main dir=outbound stage=3 role=responder result=DONE
```

```
2010-01-11 16:39:58 log_id=0101037129 type=event subtype=ipsec pri=notice vd="root" msg="progress IPsec Phase 2" action="negotiate" rem_ip=172.20.120.151 loc_ip=172.20.120.141 rem_port=500 loc_port=500 out_intf="port1" cookies="5f6dalc0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A" xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1_0" status=success init=remote mode=quick dir=outbound stage=1 role=responder result=OK
```

```
2010-01-11 16:39:58 log_id=0101037133 type=event subtype=ipsec pri=notice vd="root" msg="install IPsec SA" action="install_sa" rem_ip=172.20.120.151 loc_ip=172.20.120.141 rem_port=500 loc_port=500 out_intf="port1" cookies="5f6dalc0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A" xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1_0" role=responder in_spi=61100fe2 out_spi=bd70fca1
```

```
2010-01-11 16:39:58 log_id=0101037139 type=event subtype=ipsec pri=notice vd="root" msg="IPsec Phase 2 status change" action="phase2-up" rem_ip=172.20.120.151 loc_ip=172.20.120.141 rem_port=500 loc_port=500 out_intf="port1" cookies="5f6dalc0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A" xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1_0" phase2_name=dialup_p2
```

```
2010-01-11 16:39:58 log_id=0101037138 type=event subtype=ipsec pri=notice vd="root" msg="IPsec connection status change" action="tunnel-up" rem_ip=172.20.120.151 loc_ip=172.20.120.141 rem_port=500 loc_port=500 out_intf="port1" cookies="5f6dalc0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A" xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1_0" tunnel_ip=172.20.120.151 tunnel_id=1552003005 tunnel_type=ipsec duration=0 sent=0 rcvd=0 next_stat=0 tunnel=dialup_p1_0
```

```
2010-01-11 16:39:58 log_id=0101037129 type=event subtype=ipsec pri=notice vd="root" msg="progress IPsec Phase 2" action="negotiate" rem_ip=172.20.120.151 loc_ip=172.20.120.141 rem_port=500 loc_port=500 out_intf="port1" cookies="5f6dalc0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A" xauth_user="N/A" xauth_group="N/A" vpn_tunnel="dialup_p1_0" status=success init=remote mode=quick dir=outbound stage=1 role=responder result=OK
```

```

group="N/A" vpn_tunnel="dialup_pl_0" status=success init=remote mode=quick dir=inbound stage=2
role=responder result=DONE

2010-01-11 16:39:58 log_id=0101037122 type=event subtype=ipsec pri=notice vd="root" msg="negotiate IPsec
Phase 2" action="negotiate" rem_ip=172.20.120.151 loc_ip=172.20.120.141 rem_port=500 loc_port=500 out_
intf="port1" cookies="5f6dalc0e4bbf680/d6a1009eb1dde780" user="N/A" group="N/A" xauth_user="N/A" xauth_
group="N/A" vpn_tunnel="dialup_pl_0" status=success role=responder esp_transform=ESP_3DES esp_auth=HMAC_
SHA1

2010-01-11 16:39:58 log_id=0103031008 type=event subtype=ppp vd=root pri=information action=connect
status=success msg="Client 172.20.120.151 control connection started (id 805), assigned ip 192.168.0.50"

2010-01-11 16:39:58 log_id=0103029013 type=event subtype=ppp vd=root pri=notice pppd is started

2010-01-11 16:39:58 log_id=0103029002 type=event subtype=ppp vd=root pri=notice user="user1"
local=172.20.120.141 remote=172.20.120.151 assigned=192.168.0.50 action=auth_success msg="User 'user1'
using l2tp with authentication protocol MSCHAP_V2, succeeded"

2010-01-11 16:39:58 log_id=0103031101 type=event subtype=ppp vd=root pri=information action=tunnel-up
tunnel_id=1645784497 tunnel_type=l2tp remote_ip=172.20.120.151 tunnel_ip=192.168.0.50 user="user1"
group="L2TPUsers" msg="L2TP tunnel established"

```

Troubleshooting GRE over IPsec

This section describes some checks and tools you can use to resolve issues with the GRE-over-IPsec VPN.

Quick checks

Here is a list of common problems and what to verify.

Problem	What to check
No communication with remote network.	<p>Use the <code>execute ping</code> command to ping the Cisco device public interface.</p> <p>Use the FortiGate VPN Monitor page to see whether the IPsec tunnel is up or can be brought up.</p>
IPsec tunnel does not come up.	<p>Check the logs to determine whether the failure is in Phase 1 or Phase 2.</p> <p>Check that the encryption and authentication settings match those on the Cisco device.</p> <p>Check the encapsulation setting: tunnel-mode or transport-mode. Both devices must use the same mode.</p>
Tunnel connects, but there is no communication.	<p>Check the security policies. See Troubleshooting GRE over IPsec on page 1841.</p> <p>Check routing. See Troubleshooting GRE over IPsec on page 1841.</p>

Setting up logging

Configuring FortiGate logging for IPsec

1. Go to **Log & Report > Log Settings**.
2. Select the **Event Logging**.

3. Select **VPN activity event**.
4. Select **Apply**.

Viewing FortiGate logs

1. Go to **Log & Report > VPN Events**.
2. Select the log storage type.
3. Select **Refresh** to view any logged events.

GRE tunnel keepalives

In the event that each GRE tunnel endpoint has keepalive enabled, firewall policies allowing GRE are required in both directions. The policy should be configured as follows (where the IP addresses and interface names are for example purposes only):

```
config firewall policy
  edit < id >
    set srcintf "gre"
    set dstintf "port1"
    set srcaddr "1.1.1.1"
    set dstaddr "2.2.2.2"
    set action accept
    set schedule "always"
    set service "GRE"
  next
end
```

Cisco compatible keep-alive support for GRE

The FortiGate can send a GRE keepalive response to a Cisco device to detect a GRE tunnel. If it fails, it will remove any routes over the GRE interface.

Configuring keepalive query - CLI:

```
config system gre-tunnel
  edit <id>
    set keepalive-interval <value: 0-32767>
    set keepalive-failtimes <value: 1-255>
  next
end
```

GRE tunnel with multicast traffic

If you want multicast traffic to traverse the GRE tunnel, you need to configure a multicast policy as well as enable multicast forwarding.

- To configure a multicast policy, use the `config firewall multicast-policy` command.
- To enable multicast forwarding, use the following commands:

```
config system settings
  set multicast-forward enable
end
```

Using diagnostic commands

There are some diagnostic commands that can provide useful information. When using diagnostic commands, it is best practice that you connect to the CLI using a terminal program, such as puTTY, that allows you to save output to a file. This will allow you to review the data later on at your own speed without worry about missed data as the diag output scrolls by.

Using the packet sniffer - CLI:

1. Enter the following CLI command:

```
diag sniff packet any icmp 4
```

2. Ping an address on the network behind the FortiGate unit from the network behind the Cisco router.

The output will show packets coming in from the GRE interface going out of the interface that connects to the protected network (LAN) and vice versa. For example:

```
114.124303 gre1 in 10.0.1.2 -> 10.11.101.10: icmp: echo request
114.124367 port2 out 10.0.1.2 -> 10.11.101.10: icmp: echo request
114.124466 port2 in 10.11.101.10 -> 10.0.1.2: icmp: echo reply
114.124476 gre1 out 10.11.101.10 -> 10.0.1.2: icmp: echo reply
```

3. Enter CTRL-C to stop the sniffer.

Viewing debug output for IKE - CLI:

1. Enter the following CLI commands

```
diagnose debug application ike -1
diagnose debug enable
```

2. Attempt to use the VPN or set up the VPN tunnel and note the debug output.
3. Enter CTRL-C to stop the debug output.
4. Enter the following command to reset debug settings to default:

```
diagnose debug reset
```

Chapter 14 - Logging and Reporting

This FortiOS Handbook chapter contains the following sections:

[Logging and reporting overview](#) provides general information about logging. We recommend that you begin with this chapter as it contains information for both beginners and advanced users as well. It contains an explanation of log messages, files, and devices, and an overview of the Reporting functions.

[Logging and reporting for small networks](#) provides an overview of setting up a small network for logging, with a look at a possible setup with a backup solution and a customized report.

[Logging and reporting for large networks](#) provides an overview of setting up a larger, enterprise-level network, with configuration of multiple FortiGate units, multiple FortiAnalyzer units as a backup solution, and a sample procedure for creating a more intensive and broad report to suit the larger network.

[Advanced logging](#) provides a series of separate tutorials for possible tasks and procedures an advanced user may want to undertake with their FortiGate-powered network. It contains explanations of advanced backup, logging, and report solutions.

[Troubleshooting and logging](#) provides a short overview of how log messages can be used to identify and solve problems within the network, how to identify and solve logging database issues, and how to solve connection issues between FortiGate and FortiAnalyzer units.

What's new in FortiOS 6.0

The following list contains new Logging & Reporting features added in FortiOS 6.0.

Automatic synchronization of log display location

In previous versions, log display location could differ between Log & Report and FortiView, which could result in empty log screens if the two were not synchronized. Now, both log viewers automatically pick the best available log device. A different log device can be manually selected.

As a result, the associated CLI command `log gui-display location` has been removed.

Improved log messages for SD-WAN link quality changes

FortiOS 6.0 introduces two new log messages:

- 22923: LOG_ID_EVENT_VWL_LQTY_STATUS is created when a member's link quality is changed.
- 22924: LOG_ID_EVENT_VWL_VOLUME_STATUS is used only when `load-balance-mode` is set to `measured-volume-based`. The log is created when a member starts or stops receiving traffic.

Extended UTM logging and improved syslog configuration

Multiple UTM features now have the ability to enable extended logging: WAF, Web Filtering, DLP, AntiVirus.

These new features can be enabled in the CLI:

```
config waf profile
  edit <profile name>
    set extended-log {enable | disable}
  end
config webfilter profile
  edit <profile name>
    set web-extended-log {enable | disable}
    set web-extended-all-action-log {enable | disable}
  end
config dlp sensor
  edit <sensor name>
    set dlp-extended-log {enable | disable}
  end
config antivirus profile
  edit <profile name>
    set av-extended-log {enable | disable}
  end
```

Updated reliable syslog encryption to comply with RFC 5425

In order to align with RFC 5425 (syslog on an encrypted TLS connection over TCP) and general logging security standards for syslog, reliable syslog encryption is customizable in the CLI:

```
config log syslog setting
  set enc-algorithm {high-medium | high | low | disable}
end
```

Also, syslog options for reliable logging transmission have been expanded:

```
config log syslog setting
  set mode {udp | legacy-reliable | reliable}
end
```

See the *FortiOS CLI Reference* for more information about these commands.

Improved log display consistency at high load

Previous versions could display inconsistent log data when using Drill Down charts and when navigating between different log tables (in both **Log & Report** and **FortiView**). The maximum number of records now varies based on length that logs are kept, relative to device model size. Record numbers are configurable in `config report setting`.

Log database queries used to collect **Top Sources** and **Top Destinations** data are significantly more efficient due to improved indexing speed.

Logging and reporting overview

Logging and reporting in FortiOS can help you in determining what is happening on your network, as well as informing you of certain network activity, such as detection of a virus or IPsec VPN tunnel errors. Logging and reporting go hand in hand, and can become a valuable tool for information as well as helping to show others the activity that is happening on the network.

This section explains logging and reporting features that are available in FortiOS, and how they can be used to help you manage or troubleshoot issues. This includes how the FortiGate unit records logs, what a log message is, and what the log database is.

What is logging?

Logging records the traffic passing through the FortiGate unit to your network and what action the FortiGate unit took during its scanning process of the traffic. This recorded information is called a log message.

After a log message is recorded, it is stored within a log file which is then stored on a log device. A log device is a central storage location for log messages. The FortiGate unit supports several log devices, such as FortiAnalyzer units, the FortiCloud service, and Syslog servers. A FortiGate unit's system memory and local disk can also be configured to store logs, and because of this, are also considered log devices.



You must subscribe to FortiCloud before you will be able to configure the FortiGate unit to send logs to a FortiCloud server.

When the recorded activity needs to be read in a more human way, the FortiGate unit can generate a Report. A report gathers all the log information that is needed for the report, and presents it in a graphical format, with customizable design and automatically generated charts. Reports can be used to present a graphical representation of what is going on in the network. Reports can also be generated on a FortiAnalyzer unit; if you want to generate reports on a FortiAnalyzer, see the [FortiAnalyzer Setup and Administration Guide](#) to help you create and generate those reports.

How the FortiGate unit records log messages

The FortiGate unit records log messages in a specific order, storing them on a log device. The order of how the FortiGate unit records log messages is as follows:

1. Incoming traffic is scanned.
2. During the scanning process, the FortiGate unit performs necessary actions, and simultaneously records the actions and results.
3. Log messages are sent to the log device.

Example: How the FortiGate unit records a DLP event

1. The FortiGate unit receives incoming traffic and scans for any matches associated within its firewall policies containing a DLP sensor.
2. A match is found; the DLP sensor, `dlp_sensor`, had a rule within it called All-HTTP with the action Exempt applied to the rule. The sensor also has Enable Logging selected, which indicates to the FortiGate unit that the activity should be recorded and placed in the DLP log file.

3. The FortiGate unit exempts the match, and places the recorded activity (the log message) within the DLP log file.
4. According to the log settings that were configured, logs are stored on the FortiGate unit's local hard drive. The FortiGate unit places the DLP log file on the local hard drive.

FortiOS features available for logging

Logs record FortiGate activity, providing detailed information about what is happening on your network. This recorded activity is found in log files, which are stored on a log device. However, logging FortiGate activity requires configuring certain settings so that the FortiGate unit can record the activity. These settings are often referred to as log settings, and are found in most security profiles, but also in **Log & Report > Log Settings**.

Log settings provide the information that the FortiGate unit needs so that it knows what activities to record. This topic explains what activity each log file records, as well as additional information about the log file, which will help you determine what FortiGate activity the FortiGate unit should record.

Traffic

Traffic logs record the traffic that is flowing through your FortiGate unit. Since traffic needs firewall policies to properly flow through the unit, this type of logging is also referred to as firewall policy logging. Firewall policies control all traffic that attempts to pass through the FortiGate unit, between FortiGate interfaces, zones and VLAN sub-interfaces.

Logging traffic works in the following way:

- firewall policy has logging enabled on it (Log Allowed Traffic)
- packet comes into an inbound interface
- a possible log packet is sent regarding a match in the firewall policy, such as a URL filter
- traffic log packet is sent, per firewall policy
- packet passes and is sent out an interface

Traffic log messages are stored in the traffic log file. Traffic logs can be stored any log device, even system memory.

All security profile-related logs are now tracked within the Traffic logs, as of FortiOS 5.0, so all forward traffic can be searched in one place, such as if you are looking to see all activity from a particular address, security feature or traffic. Security profile logs are still tracked separately in the **Security Log** section, which only appears when logs exist.

If you have enabled and configured WAN Optimization, you can enable logging of this activity in the CLI using the `config wanopt setting` command. These logs contain information about WAN Optimization activity and are found in the traffic log file. When configuring logging of this activity, you must also enable logging within the security policy itself, so that the activity is properly recorded.

Sniffer

The Sniffer log records all traffic that passes through a particular interface that has been configured to act as a One-Armed Sniffer, so it can be examined separately from the rest of the Traffic logs.

Other traffic

The traffic log also records interface traffic logging, which is referred to as Other Traffic. Other Traffic is enabled only in the CLI. When enabled, the FortiGate unit records traffic activity on interfaces as well as firewall policies.

Logging Other Traffic puts a significant system load on the FortiGate unit and should be used only when necessary.

Logging other traffic works in the following way:

- firewall policy has logging enabled on it (Log Allowed Traffic) and other-traffic
- packet comes into an interface
- interface log packet is sent to the traffic log that is enabled on that particular interface
- possible log packet is sent regarding a match in the firewall policy, such as URL filter
- interface log packet is sent to the traffic log if enabled on that particular interface
- packet passes and is sent out an interface
- interface log packet is sent to traffic (if enabled) on that particular interface

Event

The event log records administration management as well as FortiGate system activity, such as when a configuration has changed, admin login, or high availability (HA) events occur. Event logs are an important log file to record because they record FortiGate system activity, which provides valuable information about how your FortiGate unit is performing.

Event logs help you in the following ways:

- keeping track of configuration setting changes
- IPsec negotiation, SSL VPN and tunnel activity
- quarantine events, such as banned users
- system performance
- HA events and alerts
- firewall authentication events
- wireless events on models with WiFi capabilities
- activities concerning modem and internet protocols L2TP, PPP and PPPoE
- VIP activities
- AMC disk's bypass mode
- VoIP activities that include SIP and SCCP protocols.

As of 5.4, every 'execute' CLI command now generates an 'audit' event log, allowing you to track configuration changes. You can enable/disable this feature in the CLI:

```
config system global
    set cli-audit-log [enable|disable]
end
```

The FortiGate unit records event logs only when events are enabled.

Traffic shaping

Traffic shaping, per-IP traffic shaping and reverse direction traffic shaping settings can be applied to a firewall policy, appearing within the traffic log messages.

By enabling this feature, you can see what traffic shaping, per-IP traffic shaping and reverse direction traffic shaping settings are being used.

Data Leak Prevention

Data Leak Prevention logs, or DLP logs, provide valuable information about the sensitive data trying to get through to your network as well as any unwanted data trying to get into your network. The DLP rules within a DLP sensor can log the following traffic types:

- email (SMTP, POP3 or IMAP; if SSL content SMTPS, POP3S, and IMAPS)
- HTTP
- HTTPS
- FTP
- NNTP
- IM

A DLP sensor must have log settings enabled for each DLP rule and compound rule, as well as applied to a firewall policy so that the FortiGate unit records this type of activity. A DLP sensor can also contain archiving options, which these logs are then archived to the log device.

NAC Quarantine

Within the DLP sensor, there is an option for enabling NAC Quarantine. The NAC Quarantine option allows the FortiGate unit to record details of DLP operation that involve the ban and quarantine actions, and sends these to the event log file. The NAC Quarantine option must also be enabled within the Event Log settings. When enabling NAC quarantine within a DLP Sensor, you must enable this in the CLI because it is a CLI-only command.

Media Access Control (MAC) address

MAC address logs provide information about MAC addresses that the FortiGate unit sees on the network as well as those removed from the network. These log messages are stored in the event log (as subtype network; you can view these log messages in **Log & Report > System Events**) and are, by default, disabled in the CLI. You can enable logging MAC addresses using the following command syntax:

```
config log setting
    set neighbor-event enable
end
```

When enabled, a new log message is recorded every time a MAC address entry is added to the ARP table, and also when a MAC address is removed as well. A MAC address log message is also recorded when MAC addresses are connected to the local switch, or from a FortiAP or FortiSwitch unit.

Application control

Application control logs provide detailed information about the traffic that internet applications such as Skype are generating. The application control feature controls the flow of traffic from a specific application, and the FortiGate unit examines this traffic for signatures that the application generates.

The log messages that are recorded provide information such as the type of application being used (such as P2P software), and what type of action the FortiGate unit took. These log messages can also help you to determine the top ten applications that are being used on your network. This feature is called application control monitoring and you can view the information from a widget on the Executive Summary page.

The application control list that is used must have logging enabled within the list, as well as logging enabled within each application entry. Each application entry can also have packet logging enabled. Packet logging for application control records the packet when an application type is identified, similar to IPS packet logging.

Logging of application control activity can only be recorded when an application control list is applied to a firewall policy, regardless of whether or not logging is enabled within the application control list.

Antivirus

Antivirus logs are recorded when, during the antivirus scanning process, the FortiGate unit finds a match within the antivirus profile, which includes the presence of a virus or grayware signature. Antivirus logs provide a way to understand what viruses are trying to get in, as well as additional information about the virus itself, without having to go to the FortiGuard Center and do a search for the detected virus. The link is provided within the log message itself.

These logs provide valuable information such as:

- the name of the detected virus
- the name of the oversized file or infected file
- the action the FortiGate unit took, for example, a file was blocked
- URL link to the FortiGuard Center which gives detailed information about the virus itself

The antivirus profile must have log settings enabled within it so that the FortiGate unit can record this activity, as well as having the antivirus profile applied to a firewall policy.

Web filter

Web filter logs record HTTP traffic activity. These log messages provide valuable and detailed information about this particular traffic activity on your network. Web filtering activity is important to log because it can inform you about:

- what types of web sites employees are accessing
- users attempting to access banned web sites and how often this occurs
- network congestion due to employees accessing the Internet at the same time
- web-based threats resulting from users visiting non-business-related web sites

Web Filter logs are an effective tool to help you determine if you need to update your web filtering settings within a web filter profile due to unforeseen threats or network congestion. These logs also inform you about web filtering quotas that have been configured for filtering HTTP traffic.

You must configure logging settings within the web filter profile and apply the filter to a firewall policy so that the FortiGate unit can record the activity.

IPS (attack)

IPS logs, also referred to as attack logs, record attacks that occurred against your network. Attack logs contain detailed information about whether the FortiGate unit protected the network using anomaly-based defense settings or signature-based defense settings, as well as what the attack was.

The IPS or attack log file is especially useful because the log messages that are recorded contain a link to the FortiGuard Center, where you can find more information about the attack. This is similar to antivirus logs, where a link to the FortiGuard Center is provided as well that informs you of the virus that was detected by the FortiGate unit.

An IPS sensor with log settings enabled must be applied to a firewall policy so that the FortiGate unit can record the activity.

Packet logs

When you enable packet logging within an IPS signature override or filter, the FortiGate unit examines network packets, and if a match is found, saves them to the attack log. Packet logging is designed to be used as a diagnostic tool that can focus on a narrow scope of diagnostics, rather than a log that informs you of what is occurring on your network.

You should use caution when enabling packet logging, especially within IPS filters. Filter configuration that contains thousands of signatures could potentially cause a flood of saved packets, which would take up a lot of storage space on the log device. It would also take a great deal of time to sort through all the log messages, as well as consume considerable system resources to process.

You can archive packets, but you must enable this option on the Log Settings page. If your log configuration includes multiple FortiAnalyzer units, packet logs are only sent to the primary (first) FortiAnalyzer unit. Sending packet logs to the other FortiAnalyzer units is not supported.

Email filter

Email filter logs, also referred to as spam filter logs, record information regarding the content within email messages. For example, within an email filter profile, a match is found that finds the email message to be considered spam.

Email filter logs are recorded when the FortiGate unit finds a match within the email filter profile and logging settings are enabled within the profile.



If you are using a Banned Words List for email filtering, note that the filter pattern number is only recorded when the source email address contains a banned word.

Archives (DLP)

Recording DLP logs for network use is called DLP archiving. The DLP engine examines email, FTP, IM, NNTP, and web traffic. Archived logs are usually saved for historical use and can be accessed at any time. IPS packet logs can also be archived, within the Log Settings page.

You can start with the two default DLP sensors that have been configured specifically for archiving log data, Content_Archive and Content_Summary. They are available in **Security Profiles > Data Leak Prevention**. Content_Archive provides full content archiving, while Content_Summary provides summary archiving. For more information about how to configure DLP sensors, see the Security Features chapter of the FortiOS Handbook.

You must enable the archiving to record log archives. Logs are not archived unless enabled, regardless of whether or not the DLP sensor for archiving is applied to the firewall policy.

Network scan

Network scan logs are recorded when a scheduled scan of the network occurs. These log messages provide detailed information about the network's vulnerabilities regarding software, as well as the discovery of any further vulnerabilities.

A scheduled scan must be configured and logging enabled within the Event Log settings, for the FortiGate unit to record these log messages.

Log messages

Log messages are recorded by the FortiGate unit, giving you detailed information about the network activity. Each log message has a unique number that helps identify it, as well as containing fields; these fields, often called log fields, organize the information so that it can be easily extracted for reports.

These log fields are organized in such a way that they form two groups: the first group, made up of the log fields that come first, is called the log header. The log header contains general information, such as the unique log identification and date and time that indicates when the activity was recorded. The log body is the second group, and contains all the other information about the activity. There are no two log message bodies that are alike, however, there may be fields common to most log bodies, such as the `srcintf` or `identidix` log fields.

The log header also contains information about the log priority level which is indicated in the `level` field. The priority level indicates the immediacy and the possible repercussions of the logged action. For example, if the field contains 'alert', you need to take immediate action with regards to what occurred. There are six log priority levels.

The log severity level is the level at and above which the FortiGate unit records logs. The log severity level is defined by you when configuring the logging location. The FortiGate unit will log all messages at and above the priority level you select. For example, if you select Error, the unit will log only Error, Critical, Alert, and Emergency level messages.

Log priority levels

Levels	Description
0 - Emergency	The system has become unstable.
1 - Alert	Immediate action is required.
2 - Critical	Functionality is affected.
3 - Error	An error condition exists and functionality could be affected.
4 - Warning	Functionality could be affected.
5 - Notification	Information about normal events.
6 - Information	General information about system operations.

The Debug priority level, not shown above, is rarely used. It is the lowest log priority level and usually contains some firmware status information that is useful when the FortiGate unit is not functioning properly.

Example log header fields

Log header	
date=(2010-08-03)	The year, month and day of when the event occurred in yyyy-mm-dd format.

Log header	
time=(12:55:06)	The hour, minute and second of when the event occurred in the format hh:mm:ss.
log_id=(2457752353)	A five or ten-digit unique identification number. The number represents that log message and is unique to that log message. This ten-digit number helps to identify the log message.
type=(dlp)	The section of system where the event occurred.
subtype=(dlp)	The subtype category of the log message.
level=(notice)	The priority level of the event. See the table above.
vd=(root)	The name of the virtual domain where the action/event occurred in. If no virtual domains exist, this field always contains root.

Example log body fields

Log body	
policyid=(1)	The ID number of the firewall policy that applies to the session or packet. Any policy that is automatically added by the FortiGate will have an index number of zero.
identidx=(0)	The identity-based policy identification number. This field displays zero if the firewall policy does not use an identity-based policy; otherwise, it displays the number of the identity-based policy entry that the traffic matched. This number is not globally unique, it is only locally unique within a given firewall policy.
sessionid=(311)	The serial number of the firewall session of which the event happened.
srcip=(10.10.10.1)	The source IP address.
srcport=(1190)	The source port number.
srcintf=(internal)	The source interface name.
dstip=(192.168.1.122)	The destination IP address.
dstport=(80)	The destination port number.
dstintf=(wan1)	The destination interface name.
service=(https)	The IP network service that applies to the session or packet. The services displayed correspond to the services configured in the firewall policy.
status=(detected)	The action the FortiGate unit took.

Log body	
hostname=(example.com)	The home page of the web site.
url=(/image/trees_pine_forest/)	The URL address of the web page that the user was viewing.
msg=(data leak detected (Data Leak Prevention Rule matched))	Explains the FortiGate activity that was recorded. In this example, the data leak that was detected matched the rule, All-HTTP, in the DLP sensor.
rulename=(All-HTTP)	The name of the DLP rule within the DLP sensor.
action=(log-only)	The action that was specified within the rule. In some rules within sensors, you can specify content archiving. If no action type is specified, this field display log-only.
severity=(1)	The level of severity for that specific rule.

Logs from other devices, such as the FortiAnalyzer unit and Syslog server, contain a slightly different log header. For example, when viewing FortiGate log messages on the FortiAnalyzer unit, the log header contains the following log fields when viewed in the Raw format:

```
itime=1302788921 date=20110401 time=09:04:23 devname=FG50BH3G09601792 device_
id=FG50BH3G09601792 log_id=0100022901 type=event subtype=system level=notice vd=root
```

The log body contains the rest of the information of the log message, and this information is unique to the log message itself.

For detailed information on all log messages, see the *FortiGate Log Message Reference*.

Explanation of a debug log message

Debug log messages are only generated if the log severity level is set to Debug. The Debug severity level is the lowest log severity level and is rarely used. This severity level usually contains some firmware status information that is useful when the FortiGate unit is not functioning properly. Debug log messages are generated by all types of FortiGate features.

The following is an example of a debug log message:

```
date=2010-01-25 time=17:25:54 logid=9300000000 type=webfilter subtype=urlfilter
level=debug msg="found in cache"
```

Example of a Debug log message

Debug log	
date=(2010-01-25)	The year, month and day of when the event occurred in the format yyyy-mm-dd.
time=(17:25:54)	The hour, minute and second of when the event occurred in the format hh:mm:ss.

Debug log	
logid=(930000000000)	A ten-digit unique identification number. The number represents that log message and is unique to that log message. This ten-digit number helps to identify the log message.
type=(webfilter)	The section of system where the event occurred. There are eleven log types in FortiOS 4.0.
subtype=(urlfilter)	The subtype of the log message. This represents a policy applied to the FortiGate feature in the firewall policy.
level=(debug)	The priority level of the event. There are six priority levels to specify.
msg=("found in cache")	Explains the activity or event that the FortiGate unit recorded.

Viewing log messages and archives

Depending on the log device, you may be able to view logs within the web-based manager or CLI on the FortiGate unit. If you have configured a FortiAnalyzer unit, local hard disk, or system memory, you can view log messages from within the web-based manager or CLI. If you have configured either a Syslog or WebTrends server, you will not be able to view log messages from the web-based manager or CLI. There is also no support for viewing log messages stored on a FortiCloud server, from the FortiGate unit's web-based manager or CLI.

You do not have to view log messages from only the web-based manager. You can view log messages from the CLI as well, using the `execute log display` command. This command allows you to see specific log messages that you already configured within the `execute log filter` command. The `execute log filter` command configures what log messages you will see, how many log messages you can view at one time (a maximum of 1000 lines of log messages), and the type of log messages you can view. For more information about viewing log messages in the CLI, see "Viewing logs from the CLI".

There are two log viewing options in FortiOS: Format and Raw. The Raw format displays logs as they appear within the log file. You can view log messages in the Raw format using the CLI or a text editor, such as Notepad. Format is in a more human-readable format, and you can easily filter information when viewing log messages this way. The Format view is what you see when viewing logs in the web-based manager.

When you download the log messages from within the log message page (for example, **Log & Report > Forward Traffic**), you are downloading log messages in the Raw format.

Viewing log messages in detail

From any log page, you can view detailed information about the log message in the log viewer table, located (by default) at the bottom of the page. Each page contains this log viewer table. The Log Viewer Table can contain the Archive tab, which allows you to see the archived version of the log message. The Archive tab only displays the archived log's details if archiving is enabled and logs are being archived by the FortiGate unit, but archived logs will also be recorded when using a FortiAnalyzer unit or the FortiCloud service.

When you are viewing traffic log messages, some of the categories (such as 'Application Name') have entries that can be selected to open a dialog box containing FortiGuard information about the entry. From within the dialog box, you can select the Reference link and go directly to the corresponding FortiGuard page, which contains additional information.

Viewing logs in Raw format allows you to view all log fields at once, as well as have a log file available regardless of whether you are archiving logs or not. You download the log file by selecting **Download Log**. The log file is named in the following format: <log_type><log_location><log_date/time>.<log_number>.log. For example, SystemEventLog-disk-2012-09-19T12_13_46.933949.log, which is an event log. The time period is the day and month of when the log was downloaded, not the time period of the log messages within the file itself.

Quarantine

Within the Log & Report menu, you can view detailed information about each quarantined file. The information can either be sorted or filtered, depending on what you want to view.

You must enable quarantine settings within an antivirus profile and the destination must be configured in the CLI using the `config antivirus quarantine` command. The destination can be either a FortiAnalyzer unit or local disk.

Sort the files by file name, date, service, status, duplicate count (DC), or time to live (TTL). Filter the list to view only quarantined files with a specific status or from a specific service.

The file quarantine list displays the following information about each quarantined file.

Quarantine page

Lists all files that are considered quarantined by the unit. On this page you can filter information so that only specific files are displayed on the page.

GUI Item	Description
Source	Either FortiAnalyzer or Local Disk , depending where you configure to quarantined files to be stored.
Sort by	Sort the list. Choose from: Status , Service , File Name , Date , TTL , or Duplicate Count . Select Apply to complete the sort.
Filter	<p>Filter the list. Choose either Status (infected, blocked, or heuristics) or Service (IMAP, POP3, SMTP, FTP, HTTP, MM1, MM3, MM4, MM7, IM, or NNTP). Select Apply to complete the filtering. Heuristics mode is configurable through the CLI only.</p> <p>If your unit supports SSL content scanning and inspection Service can also be IMAPS, POP3S, SMTPS, or HTTPS. For more information, see the Security Features chapter of the FortiOS Handbook.</p>
Apply	Select to apply the sorting and filtering selections to the list of quarantined files.
Delete	Select to delete the selected files.
Page Controls	Use the controls to page through the list.

GUI Item	Description
Remove All Entries	Removes all quarantined files from the local hard disk. This icon only appears when the files are quarantined to the hard disk.
File Name	The file name of the quarantined file. When a file is quarantined, all spaces are removed from the file name, and a 32-bit checksum is performed on the file. The checksum appears in the replacement message but not in the quarantined file. The file is stored on the FortiGate hard disk with the following naming convention: <32bit_CRC>.<processed_filename> For example, a file named Over Size.exe is stored as 3fc155d2.oversize.exe.
Date	The date and time the file was quarantined, in the format dd/mm/yyyy hh:mm. This value indicates the time that the first file was quarantined if duplicates are quarantined.
Service	The service from which the file was quarantined (HTTP, FTP, IMAP, POP3, SMTP, MM1, MM3, MM4, MM7, IM, NNTP, IMAPS, POP3S, SMTPS, or HTTPS).
Status	The reason the file was quarantined: infected , heuristics , or blocked .
Status Description	Specific information related to the status, for example, "File is infected with "W32/Klez.h"" or "File was stopped by file block pattern."
DC	Duplicate count. A count of how many duplicates of the same file were quarantined. A rapidly increasing number can indicate a virus outbreak.
TTL	Time to live in the format hh:mm. When the TTL elapses, the FortiGate unit labels the file as EXP under the TTL heading. In the case of duplicate files, each duplicate found refreshes the TTL. The TTL information is not available if the files are quarantined on a FortiAnalyzer unit.
Upload status	Y indicates the file has been uploaded to Fortinet for analysis, N indicates the file has not been uploaded. This option is available only if the FortiGate unit has a local hard disk.
Download	Select to download the corresponding file in its original format. This option is available only if the FortiGate unit has a local hard disk.
Submit	Select to upload a suspicious file to Fortinet for analysis. This option is available only if the FortiGate unit has a local hard disk.

Customizing the display of log messages on the web-based manager

Customizing log messages on the web-based manager allows you to remove or add columns from the page and filter the information that appears. For example, you can view only log messages that appeared on December 4, between the hours of 8:00 and 8:30 am.

1. Select the submenu in **Log & Report** in which you want to customize the display of log messages, such as **Log & Report > Forward Traffic**.
2. Right click on the title bar at the top of any column, and uncheck a column title such as **Date/Time** to remove it from the interface. Check other columns to add them to the interface. When you are finished, click outside the menu and the page will refresh with the new column settings in place.
3. Choose a column you'd like to filter, and select the funnel icon next to the title of the column. For example, select the funnel in the Src (Source) column. In the text field, enter the source IP address 1.1.1.1 and then select the check box beside **NOT**.
This filters out the all log messages that have the 1.1.1.1 source IP address in the source IP log field, such as the ones generated when running log tests in the CLI.
4. Select **OK** to save the customize settings, and then view the log messages on the page.
Log messages that originate from the 1.1.1.1 source address will no longer appear in the list.

How to download log messages and view them from on a computer

After recording some activity, you can download log messages to view them from a computer. This is can be very useful when in a remote location, or if you want to view log messages at your convenience, or to view packet logs or traffic logs.

1. In Log & Report, select the submenu that you want to download log messages from.
For example, **Log & Report > Forward Traffic**.
2. Select the **Download Log** option and save the log file to your computer.
The log file will be downloaded like any other file. Log file names contain their log type and date in the name, so it is recommended to create a folder in which to archive your log messages, as they can be sorted easily.
3. Open a text editor such as Notepad, open the log file, and then scroll to view all the log messages.
You can easily search or scroll through the logs to see the information that is available.

Log files and types

As the log messages are being recorded, log messages are also being put into different log files. The log file contains the log messages that belong to that log type, for example, traffic log messages are put in the traffic log file.

When downloading the log file from within **Log & Report**, the file name indicates the log type and the device on which it is stored, as well as the date, time, and a unique id for that log.

This name is in the format <logtype> - <logdevice> - <date> T <time> . <id>.log.

For example, AntiVirusLog-disk-2012-09-13T11_07_57.922495.log.

Below, each of the different log files are explained. Traffic and Event logs come in multiple types, but all contain the base type such as 'Event' in the filename.

Log Types based on network traffic

Log Type	Description
Traffic	The traffic logs records all traffic to and through the FortiGate interface. Different categories monitor different kinds of traffic, whether it be forward, local, or sniffer.
Event	The event logs record management and activity events within the device in particular areas: System, Router, VPN, User, Endpoint, HA, WAN Opt./Cache, and WiFi. For example, when an administrator logs in or logs out of the web-based manager, it is logged both in System and in User events.
Antivirus	The antivirus log records virus incidents in Web, FTP, and email traffic.
Web Filter	The web filter log records HTTP FortiGate log rating errors including web content blocking actions that the FortiGate unit performs.
Application Control	The application log records application usage, monitoring or blocking as configured in the security profiles.
Intrusion	The intrusion log records attacks that are detected and prevented by the FortiGate unit.
Email Filter	The email filter log records blocking of email address patterns and content in SMTP, IMAP, and POP3 traffic.
Vulnerability Scan	The Vulnerability Scan (Netscan) log records vulnerabilities found during the scanning of the network.
Data Leak Prevention	The Data Leak Prevention log records log data that is considered sensitive and that should not be made public. This log also records data that a company does not want entering their network.
VoIP	The VoIP log records VoIP traffic and messages. It only appears if VoIP is enabled on the Administrator Settings page.

Log database and datasets

The log database, also known as the SQL log database, is used to store logs on FortiGate units that have a built-in hard disk. The log database uses Structured Query Language (SQL), specifically it uses SQLite which is an embedded Relational Database Management System (RDBMS).



If you have disabled SQL logging and have factory defaults on the FortiGate unit, and then you upgrade the firmware, the upgrade will automatically disable SQL logging. When this occurs, you must re-enable SQL logging manually.

The FortiGate unit creates a database table for each log type, when log data is recorded. If the FortiGate unit is not recording log data, it does not create log tables for that device.

If you want to view the size of the database, as well as the log database table entries, use the `get report sql status` command. This command displays the amount of free space that is available as well as the first and last log database entry time and date.

The output of the `get report sql status` command contains information similar to the following:

```
Database size: 294912
Free size in database: 0
Database Page Size: 8192
Entry number:
Event: 49
Traffic: 370
Attack: 2
AntiVirus: 4
WebFilter: 254
AntiSpam: 2
Netscan: 18
Total: 699
First entry time: 2012-09-10 11:41:02
Last entry time: 2012-09-13 02:59:59
```

The log database is not only used to store logs, but also used to extract the information for reports. Reports are built from datasets, which are SQL statements that tell the FortiGate unit how to extract the information from the database. You can create your own datasets; however, SQL knowledge is required. Default datasets are available for reports.

Notifications about network activity

Alert email messages provide notification about activities or events logged. These email messages also provide notification about log severities that are recorded, such as a critical or emergency.

You can send alert email messages to up to three email addresses. Alert messages are also logged and can be viewed from the Event Log menu, in the System Event log file.

You can use the alert email feature to monitor logs for log messages, and to send email notification about a specific activity or event logged. For example, if you require notification about administrators logging in and out, you can configure an alert email that is sent whenever an administrator logs in and out. You can also base alert email messages on the severity levels of the logs.

Before configuring alert email, you must configure at least one DNS server if you are configuring with an Fully Qualified Domain Server (FQDN). The FortiGate unit uses the SMTP server name to connect to the mail server, and must look up this name on your DNS server. You can also specify an IP address.



The default minimum log severity level is Alert. If the FortiGate unit collects more than one log message before an interval is reached, the FortiGate unit combines the messages and sends out one alert email.

How to configure email notifications

The following explains how to configure an alert email notification for IPsec tunnel errors, firewall authentication failure, configuration changes and FortiGuard license expiry.

1. In **System > Advanced**, under **Email Service**, configure the SMTP server.

The SMTP server settings allow the FortiGate unit to know exactly where the email will be sent from, as well as who to send it to. The SMTP server must be a server that does not support SSL/TLS connections; if the SMTP server does, the alert email configuration will not work. The FortiGate unit does not currently support SSL/TLS

connections for SMTP servers.

2. In **Log & Report > Alert E-mail**, enter the source email in the Email From field, and up to three target addresses in the Email To fields.
3. Below the email entry, you can configure the email responses. By default, the **Send alert email for the following** is enabled. Select the check boxes beside **IPsec tunnel errors**, **Configuration changes** and **Firewall authentication failure**.

These alerts will be sent to the email address specified when the trigger occurs. For example, a user attempts to connect to the branch office of the company but cannot; the FortiGate unit detects an IPsec tunnel error, records the event, and then sends the notice to the email address specified in the SMTP server settings.

4. Select **FortiGuard license expiry time**: and then enter 10 so that the email notification will be sent ten days prior to the FortiGuard license expiration.
You can choose up to 100 days prior to when the license will expire. The default time is 15 days. By using this alert email notification, you can easily know when to send an re-registration request long before the expiry.

Log devices

The FortiGate unit supports a variety of log devices, including the FortiCloud service and FortiAnalyzer units. This provides greater flexibility not only when choosing a log device, but also when your logging requirements need updating.

When you have developed a plan that meets your logging needs and requirements, you need to select the log device that is appropriate for that plan. A log device must be able to store all the logs you need, and if you require archiving those logs, you must consider what log devices support this option.

During this process of deciding what log device meets your needs and requirements, you must also figure out how to provide a backup solution in the event the log device that the FortiGate unit is sending logs to has become unavailable. A backup solution should be an important part of your log setup because it helps you to maintain all logs and prevents lost logs, or logs that are not sent to the log device. For example, a daily backup of log files to the FortiAnalyzer unit occurs at 5 pm.

Log devices provide a central location for storing logs recorded by the FortiGate unit. The following are log devices that the FortiGate unit supports:

- FortiGate system memory
- Hard disk or AMC
- SQL database (for FortiGate units that have a hard disk)
- FortiAnalyzer unit
- FortiCloud service
- Syslog server

These log devices, except for the FortiGate system memory and local hard disk, can also be used as a backup solution. For example, you can configure logging to the FortiGate unit's local disk, but also configure logging to a FortiCloud server and archive logs to both the FortiCloud server and a FortiAnalyzer unit.



If you are formatting a disk that contains more than just logs, all information on the disk will be lost.

FortiGate unit's system memory and hard disk

The FortiGate unit's system memory and hard disk can store all log types, including log archives and traffic logs. Traffic logs and log archives are larger files, and need a lot of room when being logged by the FortiGate unit.

When the system memory is full, the FortiGate unit overwrites the oldest messages, and all log messages stored in memory are cleared when the FortiGate unit restarts. By default, logging to memory is enabled. This means that most of the time you will only need to modify the default settings to your network logging requirements. Real-time logging occurs whenever memory logging is enabled, and is enabled by default. Real-time logging means that the activity is being recorded as it happens.

All FortiGate units 100D and larger are capable of disk logging, but it is disabled by default, as it is not recommended. For flash memory-based units, constant rewrites to flash drives can reduce the lifetime and efficiency of the memory. For hard-disk units, it can affect performance under heavy strain. Therefore, disk logging must be manually enabled in the CLI under `config log disk setting` to appear in the interface at all.



Models without a hard disk are not recommended for disk logging. For all units, disk logging must be enabled in the CLI. For some low-end and older models, disk logging is unavailable. Check a product's Feature Matrix for more information. In either case, Fortinet recommends using either a FortiAnalyzer unit or the FortiCloud service.

Local disk or memory logging is not required for you to configure logging to a FortiAnalyzer unit.

If you are registered with the FortiCloud service, your unit will log both locally and to the service by default. In order to configure the rate and time of uploads to the service, you must register a contract account for the FortiCloud service, which will also grant you additional space.

FortiAnalyzer unit

The FortiAnalyzer unit can log all FortiGate features, which includes log archives. You can also configure the FortiGate unit to upload logs to the FortiAnalyzer unit at a scheduled time.

Encryption of the logs is supported by default and logs are sent using SSL VPN. When the FortiAnalyzer and FortiGate units have SSL encryption, both must choose a setting for the `enc-algorithm` command (CLI) for encryption to take place. By default, this is enabled and the default setting is a SSL communication with high and medium encryption algorithms. The setting that you choose must be the same for both.

FortiGate units can support logging to multiple FortiAnalyzer units. This logging solution is a backup redundancy solution, since logs are sent to all three units and whenever one of the FortiAnalyzer units fails, the others still carry on storing logs.

If you are using evaluation software FortiGate and FortiAnalyzer-VM images, you will only be able to use low-level encryption.

The FortiGate unit can also connect to a FortiAnalyzer unit using Automatic Discovery. Automatic Discovery is a method of establishing a connection to a FortiAnalyzer unit by using the FortiGate unit to find a FortiAnalyzer unit on the network. The Fortinet Discovery Protocol (FDP) is used to locate the FortiAnalyzer unit. Both the FortiGate and FortiAnalyzer units must be on the same subnet to use FDP, and they must also be able to connect using UDP.

When you enable automatic discovery in the CLI, the FortiGate unit uses HELLO packets to locate any FortiAnalyzer units that are available on the network within the same subnet. When the FortiGate unit discovers a

FortiAnalyzer unit, the FortiGate unit automatically enables logging to the FortiAnalyzer unit and begins sending log data.

Syslog server

A Syslog server is a remote computer running syslog software. Syslog is a standard for forwarding log messages in an IP network, and can be used when considering a log backup solution for your network logging requirements. Logs that are generated in real-time are sent to the syslog server in real time with no queueing, so it can be an ideal solution for comprehensive logging, or collecting logs for later systematic analysis.

FortiGate units support the reliable syslog feature, which is based on RFC 3195. Reliable syslog logging uses TCP, which ensures that connections are set up, including that packets are transmitted.

There are several profiles available for reliable syslog, but only the RAW profile is currently supported on the FortiGate units. The RAW profile is designed to provide a high-performance, low-impact footprint using essentially the same format as the existing UDP-based syslog service. The reliable syslog feature is available on FortiGate units running FortiOS 4.0 MR1 and higher.

When enabling the reliable syslog (available only in the CLI), TCP is used. The feature is disabled by default, and when enabled, the FortiGate unit automatically changes the port number to TCP 601. This is based on RFC 3195. The default port for syslog is port 514.



If you are using the local hard disk on a device for WAN Optimization, it will not prevent you from logging to remote FortiAnalyzer devices or Syslog servers. Some models have two hard disks, allowing both local logging and Wan Opt.



If you have Virtual Domains configured, each VDOM may only be assigned one FortiAnalyzer device and one Syslog server, by overriding the global configuration. The root VDOM is not limited in this way.

How to choose a log device for your network topology

When planning the log requirements, you must also consider your network's topology and whether archiving is required, such as if there is a legal requirement to keep a historical record of network activity. The following explains what steps to take when choosing a log device for your specific network topology.

1. What is the scope of your network topology?

If it is a SOHO/SMB network, then logging to the FortiGate unit's local hard disk or the default FortiCloud service would be efficient. If the network topology is a large enterprise, you will need FortiAnalyzer units, a FortiCloud contract, Syslog servers, or any combination.

2. Is archiving required?

If the network activity that is being logged needs to be archived, then, depending on your network topology, you would choose a FortiAnalyzer unit. FortiAnalyzer units store archives in the same way that FortiGate units do, but are able to store large amounts of logs and archives.

3. When troubleshooting, you may want to log a larger amount of traffic; how much storage space will you need?

Logs can be configured to roll, which is similar to zipping a file; this will lower the space requirements needed to contain them. You can also download logs from the FortiGate unit and save them on a server or on a computer to view and access later, to prevent them from piling up and being overwritten. If you're regularly logging large amounts of traffic, you should consider a FortiAnalyzer or FortiCloud account.


4. Should I invest in a log device that can grow as my network grows?

All networks grow, so investing in a device that can grow with your network and that can be expanded is a good investment. For example, if you currently have a SOHO/SMB topology, but see growth already starting, a FortiAnalyzer unit would be best. A FortiAnalyzer unit provides ample storage space, and you can add two more FortiAnalyzer units to access additional storage and create a redundancy log backup solution.

How to create a backup solution for logging

The following helps to explain how to create a log backup solution for a small network topology. This example has one FortiAnalyzer unit and a subscription to the FortiCloud Service.

Example of an integrated FortiAnalyzer unit and Syslog servers in a network



1. Log in to the CLI and modify what features will be logged to the FortiAnalyzer unit as well as the settings to the default log device, the FortiGate unit's hard drive.
By default, the FortiGate unit logs to either the system memory or hard drive, whichever is available on the FortiGate unit. Low-end FortiGate units may have logging disabled by default.
2. In the CLI, use the `config log fortianalyzer setting` command to configure logging to the FortiAnalyzer unit.
You can only configure log settings for the FortiAnalyzer unit in the CLI. Configuring to upload logs to a FortiAnalyzer unit can be configured in both the CLI and web-based manager.
3. In the CLI, configure the settings for the Syslog server; also enable reliable syslog as well.
Reliable syslog verifies that logs are sent to the syslog server. When you enable this setting, the default port becomes port 601.

Reports

Reports provide a clear, concise overview of what is happening on your network based on log data, and can be customized to serve different purposes. There are three types of reports supported by the FortiGate: FortiOS Reports, FortiCloud Reports, and FortiAnalyzer Reports.

FortiOS Reports are generated and configured on the FortiGate unit itself, FortiCloud Reports are created and configured on the FortiCloud site and mirrored to the connected FortiGate for viewing, and FortiAnalyzer reports are created and configured on a FortiAnalyzer unit. For more information about those reports, see the FortiAnalyzer Administration Guide.

In order to create FortiOS Reports on a device, disk logging must be enabled. Not all devices are capable of disk logging; check the Feature Matrix to see if your unit has a hard disk. Once disk logging has been enabled, Local Reports can then be enabled in **System > Feature Visibility** in order to view and edit reports.

What are FortiOS reports?

FortiOS reports are created from logs stored on the FortiGate unit's hard drive. These reports, generated by the FortiGate unit itself, provide a central overview of traffic and security features on the FortiGate. A default FortiOS report, called the FortiGate Security Feature Daily Activity Report, is available for you to use or modify to your requirements. The default report compiles security feature activity from various security-related logs, such as virus and attack logs. You can quickly and easily create your own report from within the management interface.

What you can do with the default FortiOS report

On the **Log & Report > Local Reports** page, you can set the frequency and timing of auto-generated reports.

You can select **Run Now** on the **Local Reports** page to immediately create a report with the current layout and design. More complex reports may take longer to generate. After generating a report, you can view it by selecting it from the list below **Run Now**.

Historical reports will be marked as 'Scheduled' if created automatically, or 'On Demand' if created by selecting **Run Now**.

What are FortiCloud reports?

FortiCloud reports are created from logs stored on the FortiCloud log management service. An active FortiCloud Service Subscription is required in order to view, configure, or use these reports. They are generated by FortiCloud according to a schedule you set, and then mirrored to the FortiGate interface and can be viewed at **Log & Report > FortiCloud Reports**, which may not appear in the interface until a report is created. If you wish to configure the report design or structure, you will have to do so from the FortiCloud portal website.

See the FortiCloud Administration Guide for more information about using and configuring FortiCloud reports.

Best practices: Log management

When the FortiGate unit records FortiGate activity, valuable information is collected that provides insight into how to better protect network traffic against attacks, including misuse and abuse. There is a lot to consider before enabling logging on a FortiGate unit, such as what FortiGate activities to enable and which log device is best suited for your network's logging needs. A plan can help you in deciding the FortiGate activities to log, a log device, as well as a backup solution in the event the log device fails.

This plan should provide you with an outline, similar to the following:

- what FortiGate activities you want and/or need logged (for example, security features)
- the logging device best suited for your network structure
- if you want or require archiving of log files
- ensuring logs are not lost in the event a failure occurs.

After the plan is implemented, you need to manage the logs and be prepared to expand on your log setup when the current logging requirements are outgrown. Good log management practices help you with these tasks.

Log management practices help you to improve and manage logging requirements. Logging is an ever-expanding tool that can seem to be a daunting task to manage. The following management practices will help you when issues arise, or your logging setup needs to be expanded.

1. Revisit your plan on a yearly basis to verify that your logging needs are being met by your current log setup. For example, your company or organization may require archival logging, but not at the beginning of your network's lifespan. Archival logs are stored on a FortiGate unit's local hard drive, a FortiAnalyzer unit, or a FortiCloud server, in increasing order of size.
2. Configure an alert message that will notify you of activities that are important to be aware about. For example: if a branch office does not have a FortiGate administrator, you will need to know at all times that the IPsec VPN tunnel is still up and running. An alert email notification message can be configured to send only if IPsec tunnel errors occur.
3. If your organization or company uses peer-to-peer programs such as Skype or other instant messaging software, use the Applications FortiView dashboard, or the Executive Summary's report widget (Top 10 Application Bandwidth Usage Per Hour Summary) to help you monitor the usage of these types of instant messaging software. These widgets can help you in determining how these applications are being used, including if there is any misuse and abuse. Their information is taken from application log messages; however, application log messages should be viewed as well since they contain the most detailed information.
4. Ensure that your backup solution is up-to-date. If you have recently expanded your log setup, you should also review your backup solution. The backup solution provides a way to ensure that all logs are not lost in the event that the log device fails or issues arise with the log device itself.

Logging and reporting for small networks

This section explains how to configure the FortiGate unit for logging and reporting in a small office or SOHO/SMB network. To properly configure this type of network, you will be modifying the default log settings, as well as the default FortiOS report.

The following procedures are examples and can be used to help you when configuring your own network's log topology. Since some of these settings must be modified or enabled or disabled in the CLI, it is recommended to review the FortiGate CLI Reference for any additional information about the commands used herein, as well as any that you would need to use in your own network's log topology.

Modifying default log device settings

The default log device settings must be modified so that system performance is not compromised. The FortiGate unit, by default, has all logging of FortiGate features enabled, except for traffic logging. The default logging location will be either the FortiGate unit's system memory or hard disk, depending on the model. Units with a flash disk are not recommended for disk logging.

Modifying the FortiGate unit's system memory default settings

When the FortiGate unit's default log device is its system memory, the following is modified for a small network topology. The following is an example of how to modify these default settings.

To modify the default system memory settings

1. Log in to the CLI.
2. Enter the following command syntax to modify the logging settings:

```
config log memory setting
    set status enable
end
```

3. The following example command syntax modifies which FortiGate features that are enabled for logging:

```
config log memory filter
    set forward-traffic enable
    set local-traffic enable
    set sniffer-traffic enable
    set anomaly enable
    set voip disable
    set multicast-traffic enable
    set dns enable
end
```

Modifying the FortiGate unit's hard disk default settings

When the FortiGate unit's default log device is its hard disk, you need to modify those settings to your network's logging needs so that you can effectively log what you want logged. The following is an example of how to modify these default settings.

To modify the default hard disk settings

1. Log in to the CLI.
2. Enter the following command syntax to modify the logging settings:

```
config log disk setting
  set ips-archive disable
  set status enable
  set max-log-file-size 1000
  set storage FLASH
  set log-quota 100
  set report-quota 100
end
```

3. In the CLI, enter the following to disable certain event log messages that you do not want logged:

```
config log eventfilter
  set event enable
  set system enable
  set vpn disable
  set user enable
  set router disable
  set wan-opt disable
end
```

Testing sending logs to the log device

After modifying both the settings and the FortiGate features for logging, you can test that the modified settings are working properly. This test is done in the CLI.

To test sending logs to the log device

1. In the CLI, enter the following command syntax:

```
diag log test
```

When you enter the command, the following appears:

```
generating a system event message with level - warning
generating an infected virus message with level - warning
generating a blocked virus message with level - warning
generating a URL block message with level - warning
generating a DLP message with level - warning
generating an IPS log message
generating an anomaly log message
generating an application control IM message with level - information
generating an IPv6 application control IM message with level - information
generating deep application control logs with level - information
generating an antispam message with level - notification
generating an allowed traffic message with level - notice
generating a multicast traffic message with level - notice
generating a ipv6 traffic message with level - notice
generating a wanopt traffic log message with level - notification
generating a HA event message with level - warning
generating netscan log messages with level - notice
generating a VOIP event message with level - information
generating a DNS event message with level - information
generating authentication event messages
generating a Forticlient message with level - information
```

```
generating a URL block message with level - warning
```

2. In the web-based interface, go to **Log & Report > System Events**, and view the logs to see some of the recently generated test log messages.
You will be able to tell the test log messages from real log messages because they do not have “real” information; for example, the test log messages for the vulnerability scan contain the destination IP address of 1.1.1.1 or 2.2.2.2.

Configuring the backup solution

A backup solution provides a way to ensure logs are not lost. The following backup solution explains logging to a FortiCloud server and uploading logs to a FortiAnalyzer unit. With this backup solution, there can be three simultaneous storage locations for logs, the first being the FortiGate unit itself, the FortiAnalyzer unit and then the FortiCloud server.

Configuring logging to a FortiCloud server

The FortiCloud server can be used as a redundant backup, or your primary logging solution. The following assumes that this service has already been registered, and a subscription has been purchased for expanded space. The following is an example of how these settings are configured for a network's log configuration. You need to have access to both the CLI and the web-based manager when configuring uploading of logs. The upload time and interval settings can be configured in the web-based interface.

To configure logging to the FortiCloud server

1. Go to **Dashboard** and click **Login** next to **FortiCloud** in the **License Information** widget.
2. Enter your username and password, and click **OK**. (Or register, if you have not yet done so.)
3. Logs will automatically be uploaded to FortiCloud as long as your FortiGate is linked to your FortiCloud account.
4. To configure the upload time and interval, go to **Log & Report > Log Settings**.
5. Under the Logging and Archiving header, you can select your desired upload time.

With FortiCloud you can easily store and access FortiGate logs that can give you valuable insight into the health and security of your network.

Configuring uploading logs to the FortiAnalyzer unit

The logs will be uploaded to the FortiAnalyzer unit at a scheduled time. The following is an example of how to upload logs to a FortiAnalyzer unit.

To upload logs to a FortiAnalyzer unit

1. Go to **Log & Report > Log Settings**.
2. In the **Remote Logging and Archiving** section, select the check box beside **Send Logs to FortiAnalyzer/FortiManager**.
3. Select **FortiAnalyzer (Daily at 00:00)**.
4. Enter the FortiAnalyzer unit's IP address in the **IP Address** field.
5. To configure the daily upload time, open the CLI.
6. Enter the following to configure when the upload occurs, and the time when the unit uploads the logs:

```
config log fortianalyzer setting
  set upload-interval {daily | weekly | monthly}
  set upload-time <hh:mm>
```

end

7. To change the upload time, in the web-based manager, select **Change** beside the upload time period, and then make the changes in the Upload Schedule window. Select **OK**.

Testing uploading logs to a FortiAnalyzer unit

You should test that the FortiGate unit can upload logs to the FortiAnalyzer unit, so that the settings are configured properly.

To test the FortiAnalyzer upload settings

1. Go to **Log & Report > Log Settings**.
2. In the **Logging and Archiving** section, under **Send Logs to FortiAnalyzer/FortiManager**, change the time to the current time by selecting **Change**.
For example, the current time is 11:10 am, so **Change** now has the time 11:10.
3. Select **OK**.

The logs will be immediately sent to the FortiAnalyzer unit, and will be available to view from within the FortiAnalyzer's interface.

Logging and reporting for large networks

This section explains how to configure the FortiGate unit for logging and reporting in a larger network, such as an enterprise network. To set up this type of network, you are modifying the default log settings, and you are also modifying the default report.

The following procedures are examples and can be used to help you when configuring your own network's log topology.

Since some of these settings must be modified or enabled or disabled in the CLI, it is recommended to review the FortiGate CLI Reference for any additional information about the commands used herein, as well as any that you would need to use in your own network's log topology.

Modifying default log device settings

The default log device settings must be modified so that system performance is not compromised. The FortiGate unit, by default, has all logging of FortiGate features enabled and well as logging to either the FortiGate unit's system memory or hard disk, depending on the model.

Modifying multiple FortiGate units' system memory default settings

When the FortiGate unit's default log device is its system memory, you can modify it to fit your log network topology. In this topic, the following is an example of how you can modify these default settings.

To modify the default system memory settings

1. Log in to the CLI.
2. Enter the following command syntax to modify the logging settings:

```
config log memory setting
    set status enable
end
```
3. Enter the following command syntax to modify the FortiGate features that are enabled for logging:

```
config log memory filter
    set forward-traffic enable
    set local-traffic enable
    set sniffer-traffic enable
    set anomaly enable
    set voip enable
    set multicast-traffic enable
    set dns enable
end
```
4. Repeat steps 2 and 3 for the other FortiGate units.
5. Test the modified settings using the procedure below.

Modifying multiple FortiGate units' hard disk default log settings

You will have to modify each FortiGate unit's hard disk default log settings. The following is an example of how to modify these default settings.

To modify the default hard disk settings

1. Log in to the CLI.
2. Enter the following command syntax to modify the logging settings:

```
config log disk setting
  set ips-archive disable
  set status enable
  set max-log-file-size 1000
  set storage Internal
  set log-quota 100
  set report-quota 100
end
```

3. In the CLI, enter the following to disable certain event log messages that you do not want logged:

```
config log eventfilter
  set event enable
  set system enable
  set vpn enable
  set user enable
  set router disable
  set wan-opt disable
end
```

4. Repeat the steps 2 to 4 for the other FortiGate units.
5. Test the modified settings using the procedure below.

Testing the modified log settings

After modifying both the settings and the FortiGate features for logging, you can test that the modified settings are working properly. This test is done in the CLI.

To test sending logs to the log device

1. In the CLI, enter the following command syntax:

```
diag log test
```

When you enter the command, the following appears:

```
generating a system event message with level - warning
generating an infected virus message with level - warning
generating a blocked virus message with level - warning
generating a URL block message with level - warning
generating a DLP message with level - warning
generating an IPS log message
generating an anomaly log message
generating an application control IM message with level - information
generating an IPv6 application control IM message with level - information
generating deep application control logs with level - information
generating an antispam message with level - notification
generating an allowed traffic message with level - notice
generating a multicast traffic message with level - notice
generating a ipv6 traffic message with level - notice
generating a wanopt traffic log message with level - notification
generating a HA event message with level - warning
generating netscan log messages with level - notice
```

```

generating a VOIP event message with level - information
generating a DNS event message with level - information
generating authentication event messages
generating a Forticlient message with level - information
generating a URL block message with level - warning

```

2. In the web-based interface, go to **Log & Report > System Events**, and view the logs to see some of the recently generated test log messages.

You will be able to tell the test log messages from real log messages because they do not have “real” information; for example, the test log messages for the vulnerability scan contain the destination IP address of 1.1.1.1 or 2.2.2.2.

Configuring the backup solution

Even though you are logging to multiple FortiAnalyzer units, this is more of a redundancy solution rather than a complete backup solution in this example.

The multiple FortiAnalyzer units act similar to a HA cluster, since if one FortiAnalyzer unit fails, the others continue storing the logs they receive. In a backup solution, the logs are backed up to another secure location if something happens to the log device.

A good alternate or redundant option is the FortiCloud service, which can provide secure online logging and management for multiple devices.

Configuring logging to multiple FortiAnalyzer units

The following example shows how to configure logging to multiple FortiAnalyzer units. Configuring multiple FortiAnalyzer units is quick and easy; however, you can only configure up to three FortiAnalyzer units per FortiGate unit.

To configure multiple FortiAnalyzer units

1. In the CLI, enter the following command syntax to configure the first FortiAnalyzer unit:

```

config log fortianalyzer setting
    set status enable
    set server 172.20.120.22
    set max-buffer-size 1000
    set buffer-max-send 2000
    set address-mode static
    set conn-timeout 100
    set monitor-keepalive-period 120
    set monitor-failure-retry-period 2000
end

```

2. Disable the features that you do not want logged, using the following example command syntax. You can view the CLI Reference to see what commands are available.

```

config log fortianalyzer filter
    set forward-traffic (enable | disable)
    ...
end

```

3. Enter the following commands for the second FortiAnalyzer unit:

```

config log fortianalyzer2 setting
    set status enable

```

```
set server 172.20.120.23
set max-buffer-size 1000
set buffer-max-send 2000
set address-mode static
set conn-timeout 100
set monitor-keepalive-period 120
set monitor-failure-retry-period 2000
end
```

4. Disable the features that you do not want logged, using the following example command syntax.

```
config log fortianalyzer2 filter
set event (enable | disable)
...
end
```

5. Enter the following commands for the last FortiAnalyzer unit:

```
config log fortianalyzer3 setting
set status enable
set server 172.20.120.23
set max-buffer-size 1000
set buffer-max-send 2000
set address-mode static
set conn-timeout 100
set monitor-keepalive-period 120
set monitor-failure-retry-period 2000
end
```

6. Disable the features that you do not want logged, using the following example command syntax.

```
config log fortianalyzer3 filter
set voip (enable | disable)
...
end
```

7. Test the configuration by using the procedure, [“Testing the modified log settings”](#).
8. On the other FortiGate units, configure steps 1 through 6, ensuring that logs are being sent to the FortiAnalyzer units.

Configuring logging to the FortiCloud server

The FortiCloud server can be used as a redundant backup, or your primary logging solution. The following assumes that this service has already been registered, and a subscription has been purchased for expanded space. The following is an example of how to these settings are configured for a network's log configuration. You need to have access to both the CLI and the web-based manager when configuring uploading of logs. The upload time and interval settings can be configured in the web-based interface.

To configure logging to the FortiCloud server

1. Go to **Dashboard** and click **Login** next to **FortiCloud** in the License Information widget.
2. Enter your username and password, and click **OK**. (Or register, if you have not yet done so.)
3. Logs will automatically be uploaded to FortiCloud as long as your FortiGate is linked to your FortiCloud account.
4. To configure the upload time and interval, go to **Log & Report > Log Settings**.
5. Under the **Remote Logging and Archiving** header, you can select your desired upload time.

6. With FortiCloud you can easily store and access FortiGate logs that can give you valuable insight into the health and security of your network.

Advanced logging

This section explains how to configure other log features within your existing log configuration. You may want to include other log features after initially configuring the log topology because the network has either outgrown the initial configuration, or you want to add additional features that will help your network's logging requirements.

The following topics are included in this section:

- [Log backup and restore tools](#)
- [Configuring logging to multiple Syslog servers](#)
- [Using Automatic Discovery to connect to a FortiAnalyzer unit](#)
- [Activating a FortiCloud account for logging purposes](#)
- [Viewing log storage space](#)
- [Customizing and filtering log messages](#)
- [Viewing logs from the CLI](#)
- [Configuring NAC Quarantine logging](#)
- [Logging local-in policies](#)
- [Tracking specific search phrases in reports](#)
- [Interpreting and configuring FSSO syslog log messages](#)

Log backup and restore tools

Local disk logs can now be backed up and restored to local files, using CLI commands:

```
execute log backup <filename>
execute log restore <filename>
```

Restoring logs will wipe the current log and report content off the disk.

Logs can also now be exported to a USB storage device, as LZ4 compressed files, from both CLI and GUI. When you insert a USB drive into the FortiGate's USB port, the USB menu will appear in the GUI. The menu shows the amount of storage on the USB disk, and the log file size, and you can select **Copy to USB** to copy the log data to the drive.

Configuring logging to multiple Syslog servers

A single remote Syslog server can be configured in the GUI, in **Log & Report > Log Settings**, but for a larger network, you will have to configure it in the CLI.

When configuring multiple Syslog servers (or one Syslog server), you can configure reliable delivery of log messages from the Syslog server. Configuring of reliable delivery is available only in the CLI.

If VDOMs are enabled, you can configure separate FortiAnalyzer unit or Syslog server for each VDOM.

To enable logging to multiple Syslog servers:

1. Log in to the CLI.
2. Enter the following commands:

```
config log syslogd setting
```

```
set csv {disable | enable}
set facility <facility_name>
set port <port_integer>
set reliable {disable | enable}
set server <ip_address>
set status {disable | enable}
end
```

3. Enter the following commands to configure the second Syslog server:

```
config log syslogd2 setting
set csv {disable | enable}
set facility <facility_name>
set port <port_integer>
set reliable {disable | enable}
set server <ip_address>
set status {disable | enable}
end
```

4. Enter the following commands to configure the third Syslog server:

```
config log syslogd3 setting
set csv {disable | enable}
set facility <facility_name>
set port <port_integer>
set reliable {disable | enable}
set server <ip_address>
set status {disable | enable}
end
```

5. Enter the following commands to configure the fourth Syslog server:

```
config log syslogd4 setting
set csv {disable | enable}
set facility <facility_name>
set port <port_integer>
set reliable {disable | enable}
set server <ip_address>
set status {disable | enable}
end
```

Most FortiGate features are, by default, enabled for logging. You can disable individual FortiGate features you do not want the Syslog server to record, as in this example:

```
config log syslogd filter
set local-traffic {enable | disable}
set severity {alert | critical | debug | emergency | error | information |
notification | warning}
end
```

Using Automatic Discovery to connect to a FortiAnalyzer unit

Automatic Discovery can be used if the FortiAnalyzer unit is on the same network.

To connect using automatic discovery

1. Log in to the CLI.
2. Enter the following command syntax:

```
config log fortianalyzer setting
set status enable
```

```
set server <ip_address>
set gui-display enable
set address-mode auto-discovery
end
```

If your FortiGate unit is in Transparent mode, the interface using the automatic discovery feature will not carry traffic. For more information about how to enable the interface to also carry traffic when using the automatic discovery feature, see the Fortinet Knowledge Base article, [Fortinet Discovery Protocol in Transparent mode](#).



The FortiGate unit searches within the same subnet for a response from any available FortiAnalyzer units.

Activating a FortiCloud account for logging purposes

When you subscribe to FortiCloud, you can configure to send logs to the FortiCloud server. The account activation can be done within the web-based manager, from the **License Information** widget located in **Dashboard**.

From this widget, you can easily create a new account, or log in to the existing account. From within the License Information widget, after the account is activated, you can go directly to the FortiCloud web portal, or log out of the service if you are already logged in.

To activate a FortiCloud account for logging purposes:

The following assumes that you are already at **Dashboard** and that you have located the License Information widget.

1. In the License Information widget, select **Activate** in the **FortiCloud** section.
The Registration window appears. From this window, you create the login credentials that you will use to access the account.
2. Select **Create Account** and enter then information for the login credentials.
After entering the login credentials, you are automatically logged in to your FortiCloud account.
3. Check that the account has been activated by viewing the account status from the License Information widget.

If you need more space, you can subscribe to the 200Gb FortiCloud service by selecting **Upgrade** in the **FortiCloud** section of the widget.

Viewing log storage space

The **Log & Report > Log Settings** GUI page displays two charts to visualize disk space: Disk Usage, which is a pie-chart illustrating the Free/Used space on the internal hard drive, and Historical Disk Usage, which displays the volume of disk logging activity over time. These charts may not be visible if disk logging is disabled.

The `diag sys logdisk usage` command allows you to view detailed information about how much space is currently being used for logs. This is useful when you see a high percentage, such as 92 percent for the disk's capacity. The FortiGate unit uses only 75 percent of the available disk capacity to avoid a high storage amount so when there is a high percentage, it refers to the percentage of the 75 percent that is available. For example, 92 percent of the 75 percent is available.

The following is an example of what you may see when you use `diag sys logdisk usage` command on a unit with no VDOMs configured:

```
diag sys logdisk usage
```

The following appears:

```
Total HD usage: 176MB/3011 MB
Total HD logging space: 22583MB
Total HD logging space for each vdom: 22583MB
HD logging space usage for vdom "root": 30MB/22583MB
```

Customizing and filtering log messages

When viewing log messages, you may want to customize and filter the information that you are seeing in the Log & Report menu (for example, **Log & Report > Forward Traffic**). Filtering and customizing the display provides a way to view specific log information without scrolling through pages of log messages to find the information.

Customizing log messages is the process of removing or adding columns to the log display page, allowing you to view certain desired information. The most columns represent the fields from within a log message, for example, the user column represents the user field, as well as additional information. If you want to reset the customized columns on the page back to their defaults, you need to select **Reset All Columns** within the column title right-click menu.

Filtering information is similar to customizing, however, filtering allows you to enter specific information that indicates what should appear on the page. For example, including only log messages that appeared on February 24, between the hours of 8:00 and 8:30 am.

To customize and filter log messages

The following is an example that displays all traffic log messages that originate from the source IP address 172.20.120.24, as well as displaying only the columns:

- OS Name
- OS Version
- Policy ID
- Src (Source IP)

The following assumes that you are already on the page of the log messages you want to customize and filter. In this example, the log messages that we are customizing and filtering are in **Log & Report > Forward Traffic**.

1. On the **Forward Traffic** page, right click anywhere on a column title.
2. Right click on a column title, and mouse over **Column Settings** to open the list.
3. Select each checkmarked title to uncheck it and remove them all from the displayed columns.
4. Scroll down to the list of unchecked fields and select 'OS Name', 'OS Version', 'Policy ID', and 'Src' to add checkmarks next to them.
5. Click outside the menu, and wait for the page to refresh with the new settings in place.
6. Select the funnel icon next to the word Src in the title bar of the Src column.
7. Enter the IP you want displayed (in this example, 172.20.120.24) in the text box.
8. Click **Apply**, and wait for the page to reload.

Viewing logs from the CLI

You can easily view log messages from within the CLI. In this example, we are viewing DLP log messages.

1. Log in to the CLI and then enter the following to configure the display of the DLP log messages.

```
execute log filter category 9
execute log filter start-line 1
```



```
execute log filter view-lines 20
```

The customized display of log messages in the CLI is similar to how you customize the display of log messages in the web-based manager. For example, `category 9` is the DLP log messages, and the `start-line` is the first line in the log database table for DLP log messages, and there will be 20 lines (`view-lines 20`) that will display.

2. Enter the following to view the log messages:

```
execute log display
```

The following appears below `execute log display`:

```
600 logs found
20 logs returned
```

along with the 20 DLP log messages.

Configuring NAC Quarantine logging

NAC Quarantine log messages provide information about what was banned and quarantined by a Antivirus profile. The following explains how to configure NAC Quarantine logging and enable it on a policy. This procedure assumes the Antivirus profile is already in place.

To configure NAC quarantine logging

1. Go to **Policy & Objects > IPv4 Policy**.
2. Select the policy that you want to apply the Antivirus profile to, and then select **Edit**.
3. Within the Security Profiles section, enable **Antivirus** and then select the profile from the drop-down list.
4. Select **OK**.
5. Log in to the CLI.
6. Enter the following to enable NAC Quarantine in the DLP sensor:

```
config antivirus profile
edit <profile_name>
config nac-quar log enable
end
```

Logging local-in policies

Local-in security policies are policies that control the flow of internal traffic, and can be used to broaden or restrict an administrator's access privileges. These local-in policies can also be configured to log traffic and activity that the policies control.

You can enable logging of local-in policies in the CLI, with the following commands:

```
config system global
set gui-local-in-policy enable
end
```

The Local-In Policy page will then be available in **Policy & Objects > Local In Policy**. You can configure what local-in traffic to log in the CLI, or in **Log & Report > Log Settings**, under **Local Traffic Logging**.

When deciding what local-in policy traffic you want logged, consider the following:

Special Traffic

Traffic activity	Traffic Direction	Description
FortiGuard update announcements	IN	All push announcements of updates that are coming from the FortiGuard system. For example, IPS or AV updates.
FortiGuard update requests	OUT	All updates that are checking for antivirus or IPS as well as other FortiGuard service updates.
Firewall authentication	IN	The authentication made using either the web-based manager or CLI.
Central management (a FortiGate unit being managed by a FortiManager unit)	IN	The access that a FortiManager has managing the FortiGate unit.
DNS	IN	All DNS traffic.
DHCP/DHCP Relay	IN	All DHCP and/or DHCP Relay traffic.
HA (heart beat sync policy)	IN/OUT	For high-end platforms with a backplane heart beat port.
HA (Session sync policy)	IN/OUT	This will get information from the CMDB and updated by session sync daemon.
CAPWAP	IN	This activity is logged only when a HAVE_CAPWAP is defined.
Radius	IN	This is recorded only within FortiCarrier.
NETBIOS forward	IN	Any interface that NETBIOS forward is enabled on.
RIP	IN	
OSPF	IN	
VRRP	IN	
BFD	IN	
IGMP	IN	This is recorded only when PIM is enabled.
PIM	IN	This is recorded only when PIM is enabled.

Traffic activity	Traffic Direction	Description
BGP	IN	This is recorded only when config bgp and bgp neighbor is enabled in the CLI.
WCCP policy	IN	Any interface that WCCP is enabled; however, if in Cache mode, this is not recorded because it is not available.
WAN Opt/ Web Cache	IN	Any interface where WAN Opt is enabled.
WANOpt Tunnel	IN	This is recorded when HAVE_WANOPT is defined.
SSL-VPN	IN	Any interface from a zone where the action in the policy is SSL VPN.
IPSEC	IN	
L2TP	IN	
PPTP	IN	
VPD	IN	This is recorded only when FortiClient is enabled.
Web cache db test facility	IN	This is recorded only when WA_CS_REMOTE_TEST is defined.
GDBserver	IN	This is recorded only when debug is enabled.

Tracking specific search phrases in reports

It is possible to use the Web Filter to track specific search keywords and phrases and record the results for display in the report.

You should verify that the web filter profile you are using indicates what search phrases you want to track and monitor, so that the report includes this information.

1. Log in to the CLI and enter show webfilter profile default.

This provides details about the webfilter profile being used by the security policy. In this example, the details (shown in the following in bold) indicate that safe search is enabled, but not specified or being logged.

```
show webfilter profile default
config webfilter profile
edit "default"
    set comment "default web filtering"
    set inspection-mode flow-based
    set options https-scan
    set post-action comfort
    config web
        set safe-search url
    end
    config ftgd-wf
        config filters
            edit 1
```

```

        set action block
        set category 2
    next
    edit 2
        set action block
        set category 7
    next
    edit 3
        set action block
        set category 8

```

2. Enter the following command syntax so that logging and the keyword for the safe search will be included in logging.

```

config webfilter profile
  edit default
    config web
      set log-search enable
      set keyword-match "fortinet" "easter" "easter bunny"
    end
  end
end

```

3. To test that the keyword search is working, go to a web browser and begin searching for the words that were included in the webfilter profile, such as easter.

You can tell that the test works by going to **Log & Report > Forward Traffic** and viewing the log messages.

Interpreting and configuring FSSO syslog log messages

There are two syslog message formats: default and verbose. Verbose must be manually enabled as described below, but provides more general information.

Default syslog message format

The default FSSO syslog message format has no header, and is based on the specifications of [RFC 3164](#). Messages only have two values, `PRI` (Priority) and `MSG` (Message), in the format of `<PRI>MSG`.

The content of `PRI` is as described in RFC 3164, but with specific parameters: the Facility value is always 1 (USER), unless 'Log logons in separate log' is enabled in the FSSO Collector Agent settings. In that case, those logon messages will have a Facility value of 4 or 10 (AUTH). The Severity value always matches the internal severity value of the log. `PRI` is enclosed in `<>` with no space following before `MSG`.

Verbose syslog message format

Verbose is a secondary message format that provides more information, including timestamp (with timezone).

In verbose mode, the log message follows the specifications of [RFC 5424](#):

```

<PRI>VERSION TIMESTAMP HOSTNAME APP-NAME PROCID MSGID STRUCTURED-DATA/SD-ID
MSG

```

`PRI` is formatted as described above in the default format.

Verbose FSSO syslog messages do not contain any data for `MSGID`, or `STRUCTURED-DATA`, so both of those two messages are recorded as a single hyphen character `"-"`.

`APP-NAME` always appears as `"collectoragent"`.

The other values are formatted as described in RFC 5424.

Enabling verbose syslog message mode

In order to enable the verbose syslog message mode, you must modify the registry on the PC that is hosting the FSSO Collector Agent.

In 64-bit Windows, locate the following registry entry:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Wow6432Node\Fortinet\FSAE\collectoragent
```

In 32-bit Windows, locate the following registry entry:

```
HKEY_LOCAL_MACHINE\SOFTWARE\Fortinet\FSAE\collectoragent
```

Under this registry path, create a new DWORD (32bit) Value named `syslog_using_rfc`, and set its value to 1.

Troubleshooting and logging

This section explains how to troubleshoot logging configuration issues, as well as connection issues, that you may have with your FortiGate unit and a log device. This section also contains information about how to use log messages when troubleshooting issues that are about other FortiGate features, such as VPN tunnel errors.

Using log messages to help in troubleshooting issues

Log messages can help when troubleshooting issues that occur, since they can provide details about what is occurring. The uses and methods for involving logging in troubleshooting vary depending on the problem. The following are examples of how log messages can assist when troubleshooting networking issues.

Using IPS packet logging in diagnostics

This type of logging should only be enabled when you need to know about specific diagnostic information, for example, when you suspect a signature is triggered by a false positive. These log messages can help troubleshoot individual problems with misidentified or missing packets and network intrusions involving malicious packets.

To configure IPS packet logging

1. Go to **Security Profiles > Intrusion Protection**.
2. Select the IPS sensor that you want to enable IPS packet logging on, and then select **Edit**.
3. In the filter options, enable **Packet Logging**.
4. Select **OK**.

If you want to configure the packet quota, number of packets that are recorded before alerts and after attacks, use the following procedure.

To configure additional settings for IPS packet logging

1. Log in to the CLI.
2. Enter the following to start configuring additional settings:

```
config ips settings
    set ips-packet-quota <integer>
    set packet-log-history <integer>
    set packet-log-post-attack <integer>
end
```

Using HA log messages to determine system status

When the FortiGate unit is in HA mode, you may see the following log message content within the event log:

```
type=event subtype=ha level=critical msg= "HA slave heartbeat interface internal lost
neighbor information"
```

OR

```
type=event subtype=ha level=critical msg= "Virtual cluster 1 of group 0 detected new
joined HA member"
```

OR

```
type=event subtype=ha level=critical msg= "HA master heartbeat interface internal get peer information"
```

The log messages occur within a given time, and indicate that the units within the cluster are not aware of each other anymore. These log messages provide the information you need to fix the problem.

Connection issues between FortiGate unit and logging devices

If external logging devices are not recording the log information properly or at all, the problem will likely be due to one of two situations: no data is being received because the log device cannot be reached, or no data is being sent because the FortiGate unit is no longer logging properly.

Unable to connect to a supported log device

After configuring logging to a supported log device, and testing the connection, you may find you cannot connect. To determine whether this is the problem:

1. Verify that the information you entered is correct; it could be a simple mistake within the IP address or you may have not selected **Apply** on the Log Settings page after changing them, which would prevent them from taking effect.
2. Use `execute ping` to see if you can ping to the log device.
3. If you are unable to ping to the log device, check to see if the log device itself working and that it is on the network and assigned an appropriate address.

FortiGate unit has stopped logging

If the FortiGate unit stopped logging to a device, test the connection between both the FortiGate unit and device using the `execute ping` command. The log device may have been turned off, is upgrading to a new firmware version, or just not working properly.

The FortiGate unit may also have a corrupted log database. When you log into the web-based manager and you see an SQL database error message, it is because the SQL database has become corrupted. View "SQL database errors" in the next section before taking any further actions, to avoid losing your current logs.

Log database issues

If attempting to troubleshoot issues with the SQL log database, use the following to help guide you to solving issues that occur.

SQL statement syntax errors

There may be errors or inconsistencies in the SQL used to maintain the database. Here are some example error messages and possible causes:

```
You have an error in your SQL syntax (remote/MySQL)
```

or

```
ERROR: syntax error at or near... (local/PostgreSQL)
```

- Verify that the SQL keywords are spelled correctly, and that the query is well-formed.
- Table and column names are demarked by grave accent (`) characters. Single (') and double (") quotation marks will cause an error.

```
No data is covered.
```

- The query is correctly formed, but no data has been logged for the log type. Verify that you have configured the FortiGate unit to save that log type. On the Log Settings page, make sure that the log type is checked.

Connection problems

If well-formed SQL queries do not produce results, and logging is turned on for the log type, there may be a database configuration problem with the remote database.

Ensure that:


- MySQL is running and using the default port 3306.
- You have created an empty database and a user who has read/write permissions for the database.
- Here is an example of creating a new MySQL database named fazlogs, and adding a user for the database:

1. `#Mysql -u root -p`
2. `mysql> Create database fazlogs;`
3. `mysql> Grant all privileges on fazlogs.* to 'fazlogger'@'*' identified by 'fazpassword';`
4. `mysql> Grant all privileges on fazlogs.* to 'fazlogger'@'localhost' identified by 'fazpassword';`

SQL database errors

If the database seems inaccessible, you may encounter the following error message after upgrading or downgrading the FortiGate unit's firmware image.

Example of an SQL database error message



The error message indicates that the SQL database is corrupted and cannot be updated with the SQL schemas any more. When you see this error message, you can do one of the following:

- select **Cancel** and back up all log files; then select **Rebuild** to blank and rebuild the database.
- select **Rebuild** immediately, which will blank the database and previous logs will be lost.

Until the database is rebuilt, no information will be logged by the FortiGate unit regardless of the log settings that are configured on the unit. When you select **Rebuild**, all logs are lost because the SQL database is erased and then rebuilt again. Logging resumes automatically according to your settings after the SQL database is rebuilt.

To view the status of the database, use the `diagnose debug sqlldb-error status` command in the CLI. This command will inform you whether the database has errors present.

If you want to view the database's errors, use the `diagnose debug sqlldb-error read` command in the CLI. This command indicates exactly what errors occurred, and what tables contain those errors.

Log files are backed up using the `execute backup {disk | memory} {alllogs | logs}` command in the CLI. You must use the text variable when backing up log files because the text variable allows you to view the log files outside the FortiGate unit. When you back up log files, you are really just copying the log files from the database to a specified location, such as a TFTP server.

Logging daemon (Miglogd)

The number of logging daemon child processes has been made available for editing. A higher number can affect performance, and a lower number can affect log processing time, although no logs will be dropped or lost if the number is decreased.

If you are suffering from performance issues, you can alter the number of logging daemon child processes, from 0 to 15, using the following syntax. The default is 8.

```
config system global
    set miglogd-children <integer>
end
```

Chapter 15 - Managing Devices

[What's new in FortiOS 6.0](#)

[Managing “bring your own device”](#)

This handbook chapter contains the following sections:

[Managing “bring your own device”](#) describes device monitoring, devices, device groups, and device policies. The administrator can monitor all types of devices and control their access to network resources.

What's new in FortiOS 6.0

The following list contains new device management features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- ["Device organization, device categories, and device types" on page 1894](#)

Managing “bring your own device”

FortiOS can control network access for different types of personal mobile devices that your employees bring onto your premises. You can:

- identify and monitor the types of devices connecting to your networks, wireless or wired
- use MAC address based access control to allow or deny individual devices
- create security policies that specify device types
- enforce endpoint control on devices that can run FortiClient Endpoint Control software









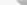
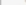



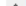


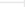












This chapter contains the following sections:

Device monitoring

The FortiGate unit can monitor your networks and gather information about the devices operating on those networks. Collected information includes:

- MAC address
- IP address
- operating system
- hostname
- user name
- how long ago the device was detected and on which FortiGate interface

You can go to **User & Device > Device Inventory** to view this information. Mouse-over the **Device** column for more details.

 Edit	 Delete	 Refresh	 Search	By Type	By Interface	Alphabetically	Total Devices Tracked: 47
 Status	 Device	 OS	 User	 IP Address	 Interface		
 Online	 00:12:7f:4d:4d:97				 wan1		
 Online	 00:14:a9:52:23:82				 wan1		
 Online	 amoffitt-pc	Windows / 7, 8 (x86)		172.20.120.51	 wan1		
 Online	 DAHLIA			172.20.121.150	 wan1		
 Offline	 00:09:0f:09:de:12			172.20.121.2	 wan1		
 Offline	 00:0c:29:07:ae:75				 wan1		

Depending on the information available, the Device column lists the Alias or the MAC address of the device. For ease in identifying devices, Fortinet recommends that you assign each device an Alias.

Device monitoring is enabled separately on each interface. Device detection is intended for devices directly connected to your LAN ports. If enabled on a WAN port, device detection may be unable to determine the operating system on some devices. Hosts whose device type cannot be determined passively can be found by enabling active scanning on the interface.

You can also manually add devices. This enables you to ensure that a device with multiple interfaces is displayed as a single device.

To configure device monitoring

1. Go to **Network > Interfaces**.
2. Edit the interface that you want to monitor devices on.
3. In **Networked Devices**, turn on **Device Detection** and optionally turn on **Active Scanning**.
4. Select **OK**.
5. Repeat steps 2 through 4 for each interface that will monitor devices.

To assign an alias to a detected device or change device information

1. Go to **User & Device > Device Inventory** and edit the device entry.
2. Enter an **Alias** such as the user's name to identify the device.
3. Change other information as needed.
4. Select **OK**.

To add a device manually

1. Go to **User & Device > Custom Devices & Groups**.
2. Select **Create New > Device**.
3. Enter the following information:
 - Alias (required)
 - MAC address
 - Additional MACs (other interfaces of this device)
 - Device Type
 - Optionally, add the device to **Custom Groups**.
 - Optionally, enter **Comments**.
3. Select **OK**.

Custom avatars for custom devices

You can upload an avatar for a custom device. The avatar is then displayed in the GUI wherever the device is listed, such as FortiView, log viewer, or policy configuration. To upload an avatar image, click **Upload Image** on the New Device or Edit Device page of **User & Device > Custom Devices & Groups**. The image can be in any format your browser supports and will be automatically sized to 36 x 36 pixels for use in the FortiGate GUI.

Device offline timeout

A device is considered offline if it has not sent any packets during the timeout period. The timeout can be set to any value from 30 to 31 536 000 seconds (365 days). The default value is 300 seconds (5 minutes). The timer is configurable in the CLI:

```
config system global
    set device-idle-timeout 300
end
```

Device organization, device categories, and device types

A second level of organization shows device category (except for device types not listed below). The categories, along with the devices that belong to those categories, include:

Category	Devices
Android	Android Phone, Android Tablet
BlackBerry	BlackBerry Phone, BlackBerry Playbook
Fortinet	Fortinet Device, FortiCam, FortiFone
iOS	iPad, iPhone
Windows	Windows PC, Windows Phone, Windows Tablet

Syntax

```
config user device
  edit <category>
    set category [none | android-device | blackberry-device | fortinet-device | ios-
device | windows-device]
  next
end
```

Device groups

You can specify multiple device types in a security policy. As an alternative, you can add multiple device types to a custom device group and include the group in the policy. This enables you to create a different policy for devices that you know than for devices in general.

To create a custom device group and add devices to it

1. Go to **User & Device > Custom Devices & Groups**.
The list of device groups is displayed.
2. Select **Create New > Device Group**.
3. Enter a **Name** for the new device group.
4. Click in the **Members** field and click a device type to add. Repeat to add other devices.
5. Select **OK**.

Controlling access with a MAC Address Access Control List

A MAC Address Access Control List (ACL) allows or blocks access on a network interface that includes a DHCP server. If the interface does not use DHCP, or if you want to limit network access to a larger group such as employee devices, it is better to create a device group and specify that group in your security policies.

A MAC Address ACL functions as either

- a list of devices to block, allowing all other devices

or

- a list of devices to allow, blocking all other devices

Allowed devices are assigned an IP address. The Assign IP action assigns the device an IP address from the DHCP range. In a list of allowed devices, you can also use the Reserve IP action to always provide a specific IP address to the device.

The **Unknown MAC Address** entry applies to "other" unknown, unlisted devices. Its action must be opposite to that of the other entries. In an allow list, it must block. In a block list, it must allow.

To create a MAC Address ACL to allow only specific devices

1. Go to the SSID or network interface configuration.
2. In the **DHCP Server** section, expand **Advanced**.
DHCP Server must be enabled.
3. In **MAC Reservation + Access Control**, select **Create New** and enter an allowed device's **MAC Address**.
4. In the **IP or Action** column, select one of:
 - Assign IP — device is assigned an IP address from the DHCP server address range.
 - Reserve IP — device is assigned the IP address that you specify.
5. Repeat Steps ["Controlling access with a MAC Address Access Control List" on page 1894](#) and ["Controlling access with a MAC Address Access Control List" on page 1894](#) for each additional MAC address entry.
6. Set the **Unknown MAC Address** entry **IP or Action** to **Block**.
Devices not in the list will be blocked.
7. Select **OK**.

To create a MAC Address ACL to block specific devices

1. Go to the SSID or network interface configuration.
2. In the **DHCP Server** section, expand **Advanced**.
DHCP Server must be enabled.
3. In **MAC Reservation + Access Control**, select **Create New** and enter the **MAC Address** of a device that must be blocked.
4. In the **IP or Action** column, select **Block**.
5. Repeat Steps ["Controlling access with a MAC Address Access Control List" on page 1894](#) and ["Controlling access with a MAC Address Access Control List" on page 1894](#) for each device that must be blocked.
6. Set the **Unknown MAC Address** entry **IP or Action** to **Assign IP**.
Devices not in the list will be assigned IP addresses.
7. Select **OK**.

MAC Authentication Bypass (MAB)

MAC Authentication Bypass allows devices without 802.1X capability (printers and IP phones for example) to bypass authentication and be allowed network access based on their MAC address. This feature requires RADIUS-based 802.1X authentication in which the RADIUS server contains a database of authorized MAC addresses.

MAC Authentication Bypass is configurable only in the CLI and only on interfaces configured for 802.1X authentication. For example:

```
config system interface
edit "lan"
set ip 10.0.0.200 255.255.255.0
```

```
        set vlanforward enable
        set security-mode 802.1X
        set security-mac-auth-bypass enable
        set security-groups "Radius-group"
    end
end
```

MAC Authentication Bypass is also available on WiFi SSIDs, regardless of authentication type. It is configurable only in the CLI. You need to enable the `radius-mac-auth` feature and specify the RADIUS server that will be used. For example:

```
config wireless-controller vap
    edit "office-ssid"
        set security wpa2-only-enterprise
        set auth usergroup
        set usergroup "staff"
        set radius-mac-auth enable
        set radius-mac-auth-server "ourRadius"
    end
end
```














Security policies for devices

Security policies enable you to implement policies according to device type. For example:














- Gaming consoles cannot connect to the company network or the Internet.
- Personal tablet and phone devices can connect to the Internet but not to company servers.
- Company-issued laptop computers can connect to the Internet and company servers. Web filtering and antivirus are applied.
- Employee laptop computers can connect to the Internet, but web filtering is applied. They can also connect to company networks, but only if FortiClient Endpoint Security is installed to protect against viruses.

The following images show these policies implemented for WiFi to the company network and to the Internet.

Device policies for company laptop access to the company network

New Policy	
Name	Laptop LAN Access
Incoming Interface	 internal (lan) 
Outgoing Interface	 lan 
Source	 1st floor LAN   employee laptop 
Destination Address	 all 
Schedule	always 
Services	 ALL 
Action	ACCEPT DENY

Device policies for WiFi access to the Internet

New Policy	
Name	WiFi access to Internet
Incoming Interface	 example-staff (example-wifi) 
Outgoing Interface	 wan1 
Source	 all   Android Phone   Android Tablet   BlackBerry Phone   BlackBerry PlayBook   iPad   iPhone 
Destination Address	 all 
Schedule	always 
Services	 ALL 
Action	ACCEPT DENY

The next section explains device policy creation in detail.

Creating device policies

Device-based security policies are similar to policies based on user identity:

- The policy enables traffic to flow from one network interface to another.
- NAT can be enabled.
- UTM protection can be applied.

To create a device policy

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Choose **Incoming Interface**, **Outgoing Interface** and **Source** as you would for any security policy.
3. In **Source**, select an address and the device types that can use this policy.
You can select multiple devices or device groups.
4. Turn on **NAT** if appropriate.
5. Configure **Security Profiles** as you would for any security policy.
6. Select **OK**.

Adding endpoint protection

Optionally, you can require that users' devices connecting to a particular network interface have FortiClient Endpoint Security software installed. Devices without an up-to-date installation of FortiClient software are restricted to a captive portal from which the user can download a FortiClient installer. For information about creating FortiClient profiles, see "Endpoint Protection".

To add endpoint protection to a security policy

1. Go to **Network > Interfaces** and edit the interface.
2. In **Admission Control** turn on **Allow FortiClient Connections** and **FortiClient Enforcement**.
3. Optionally, select sources (addresses and device types) to exempt from FortiClient enforcement.
4. Optionally, select destination addresses and services to exempt from FortiClient enforcement.
5. Select **OK**.

FortiOS pushes a FortiClient profile out to the FortiClient software, configuring network protection such as antivirus, application control, and web category filtering. To create these profiles, go to **Security Profiles > FortiClient Profiles**.

FortiClient endpoint licence updates

FortiClient endpoint licenses for FortiOS 5.6.0 can be purchased in multiples of 100. There is a maximum client limit based on the FortiGate's model. FortiCare enforces the maximum limits when the customer is applying the license to a model.

If you are using the ten free licenses for FortiClient, support is provided on the Fortinet Forum (forum.fortinet.com). Phone support is only available for paid licenses.

Model(s)	Maximum Client Limit
VM00	200
FGT/FWF 30 to 90 series	200
FGT 100 to 400 series	600
FGT 500 to 900 series, VM01, VM02	2,000
FGT 1000 to 2900 series	20,000
FGT 3000 to 3600 series, VM04	50,000
FGT 3700D and above, VM08 and above	100,000

Older FortiClient SKUs will still be valid and can be applied to FortiOS 5.4 and 5.6.

Chapter 16 - FortiSwitch Devices Managed by FortiOS

6.0.1

Introduction

NOTE: FortiLink is not supported in transparent mode.

The maximum number of supported FortiSwitch units depends on the FortiGate model:

FortiGate Model Range	Number of FortiSwitch Units Supported
Up to FortiGate-98 and FortiGate-VM01	8
FortiGate-100 to 280 and FortiGate-VM02	24
FortiGate-300 to 5xx	48
FortiGate-600 to 900 and FortiGate-VM04	64
FortiGate-1000 and up	128
FortiGate-3xxx and up and FortiGate-VM08 and up	300

Supported models

The following table shows the FortiSwitch models that support FortiLink mode when paired with the corresponding FortiGate models and the listed minimum software releases. For example, the FGT-500E model with FortiOS 5.6.3 and later supports all FortiSwitch D-series and E-series models running FortiSwitchOS 3.6.0 and later.

Each row includes support for earlier FortiGate models. For example, the FGT-500E row includes support by the FortiGate models in the rows above it.

FortiGate and FortiWiFi Models	Earliest FortiOS	FortiSwitch Models
FGT-90D	5.2.2	FS-224D-POE

FortiGate and FortiWiFi Models	Earliest FortiOS	FortiSwitch Models
FGT-60D FGT-100D, 140D, 140D-POE, 140D-T1 FGT-200D, 240D, 280D, 280D-POE FGT-600C FGT-800C FGT-1000C, 1200D, 1500D FGT-3700D, FGT-3700DX	5.2.3	FSR-112D-POE FS-108D-POE FS-124D (POE) FS-224D-POE and FPOE
	5.4.0	All FortiSwitch D-series models. FortiSwitchOS 3.3.x or 3.4.0 is recommended.
FGT and FWF-30D, 30D-POE, 30E FGT and FWF-50E, 51E FGR-60D FGT-70D, 70D-POE FGT-80D FGR-90D FGT and FWF-92D FGT-94D-POE, 98D-POE FGT-300D FGT-400D FGT-500D FGT-600D FGT-900D FGT-1000D FGT-3000D, 3100D, 3200D, 3240C, 3600C, 3810D, 3815D FGT_VM, VM64, VM64-AWS, VM64-AWSONDEMAND, VM64-HV, VM64-KVM, VM-VMX, VM64-XEN	5.4.1	All FortiSwitch D-series models. FortiSwitchOS 3.4.2 or later is required for all managed switches.
FGT and FWF- 60E, 61E FGT-100E, 101E	5.4.2	All FortiSwitch D-series models. FortiSwitchOS 3.4.2 or later is required for all managed switches.
FGT-80E, 80E-POE, 81E, 81E-POE FGT-100EF	5.4.3	All FortiSwitch D-series models. FortiSwitchOS 3.4.2 or later is required for all managed switches.
FGT-90E, 91E FGT-200E, 201E FGT-2000E, 2500E	5.6.0	All FortiSwitch D-series models. FortiSwitchOS 3.5.4 or later is required for all managed switches.

FortiGate and FortiWiFi Models	Earliest FortiOS	FortiSwitch Models
FGT-500E	5.6.3	<p>All FortiSwitch D-series and E-series models.</p> <p>FortiSwitchOS 3.6.0 or later is required for all managed switches.</p>

Support of FortiLink features

The following table lists the FortiSwitch models supported by FortiLink features.

FortiLink Features	FortiSwitch Models
Centralized VLAN Configuration	D-series, E-series
Switch POE Control	D-series, E-series
Link Aggregation Configuration	D-series, E-series
Spanning Tree Protocol (STP)	D-series, E-series
LLDP/MED	D-series, E-series
IGMP Snooping	Not supported on 112D-POE, 1xxE-Series
802.1x Authentication (Port-based, MAC-based, MAB)	D-series, E-series
Syslog Collection	D-series, E-series
DHCP Snooping	Not supported on 1xxE-Series
Device Detection	D-series, E-series
Support FortiLink FortiGate in HA Cluster	D-series, E-series
LAG support for FortiLink Connection	D-series, E-series
Active-Active Split MLAG from FortiGate to FortiSwitch units for Advanced Redundancy	Not supported on FS-1xx Series
sFlow	Not supported on 1xxE-Series
Dynamic ARP Inspection (DAI)	Not supported on 1xxE-Series
Port Mirroring	D-series, E-series
RADIUS Accounting Support	Not supported on 1xxE-Series

FortiLink Features	FortiSwitch Models
Centralized Configuration	D-series, E-series
Access VLAN	Not supported on 1xxE-Series, 112D-POE
STP BPDU Guard, Root Guard, Edge Port	D-series, E-series
Loop Guard	D-series, E-series
Switch admin Password	D-series, E-series
Storm Control	D-series, E-series
802.1x-Authenticated Dynamic VLAN Assignment	D-series, E-series
Host Quarantine on Switch Port	Not supported on 1xxE Series, 112D-POE
QoS	Not supported on 1xxE-Series, 112D-POE
Centralized Firmware Management	D-series, E-series

Before you begin

Before you configure the managed FortiSwitch unit, the following assumptions have been made in the writing of this manual:

- You have completed the initial configuration of the FortiSwitch unit, as outlined in the QuickStart Guide for your FortiSwitch model, and you have administrative access to the FortiSwitch web-based manager and CLI.
- You have installed a FortiGate unit on your network and have administrative access to the FortiGate web-based manager and CLI.

How this guide is organized

This guide contains the following sections:

- [What's new in FortiOS 6.0.1](#) describes the new features for this release.
- [Connecting FortiLink ports](#) describes how to connect FortiSwitch ports to FortiGate ports.
- [FortiLink configuration using the FortiGate GUI](#) describes how to use the FortiGate GUI for FortiLink configuration.
- [FortiLink configuration using the FortiGate CLI](#) describes how to use the FortiGate CLI for FortiLink configuration.
- [Network topologies for managed FortiSwitch units](#) describes the configuration for various network topologies.
- [Optional setup tasks](#) describes other setup tasks that are optional.
- [FortiSwitch features configuration](#) describes how to configure managed FortiSwitch features, including VLANs.
- [FortiSwitch port features](#) describe how to configure ports and PoE from the FortiGate unit.
- [FortiSwitch port security policy](#) describes how to set up FortiSwitch security policies.
- [Additional capabilities](#) describes more FortiSwitch features.
- [Troubleshooting](#) describes techniques for troubleshooting common problems.

What's new in FortiOS 6.0.1

The following list contains new managed FortiSwitch features added in FortiOS 6.0.1. Click on a link to navigate to that section for further information.

- ["Logging violations of the MAC address learning limit \(480808\)" on page 1904](#)
- ["Testing 802.1x authentication with monitor mode \(480807\)" on page 1905](#)
- ["Checking FortiSwitch connections \(453491\)" on page 1905](#)
- ["CLI changes for quarantining MAC addresses \(442831\)" on page 1906](#)
- ["CLI changes for releasing MAC addresses from quarantine \(442831\)" on page 1907](#)

What's new in FortiOS 6.0

The following list contains new managed FortiSwitch features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- ["Limiting the number of learned MAC addresses on a FortiSwitch interface \(445087\)" on page 1907](#)
- ["Sharing FortiSwitch ports between VDOMs \(391878\)" on page 1908](#)
- ["sFlow support \(450507\)" on page 1910](#)
- ["Restricting the type of frames allowed through IEEE 802.1Q ports \(448505\)" on page 1911](#)
- ["Dynamic ARP inspection \(DAI\) support \(462511\)" on page 1912](#)
- ["FortiSwitch port mirroring support \(457122\)" on page 1912](#)
- ["Quarantining MAC addresses \(459525\)" on page 1913](#)
- ["Banning IP addresses \(459525\)" on page 1913](#)
- ["Synchronizing the FortiGate unit with the managed FortiSwitch units \(454664\)" on page 1914](#)
- ["Enabling the use of HTTPS to download firmware to managed FortiSwitch units \(454664\)" on page 1914](#)
- ["RADIUS accounting support \(451023\)" on page 1914](#)
- ["FortiLink mode supported over a layer-3 network \(457103\)" on page 1915](#)
- ["Limiting the number of parallel process for FortiSwitch configuration \(457103\)" on page 1916](#)
- ["CLI changes for FortiLink mode \(447349, 473773\)" on page 1917](#)
- ["Upgrade the firmware on multiple FortiSwitch units at the same time using the GUI \(462553\)" on page 1917](#)
- ["Network-assisted device detection \(377467\)" on page 1918](#)

FortiOS 6.0.1

These features first appeared in FortiOS 6.0.1.

Logging violations of the MAC address learning limit (480808)

If you set a maximum number of MAC addresses learned for an interface or VLAN, the managed FortiSwitch unit drops all traffic for additional MAC addresses after the learning limit is reached.

You can now change how long learned MAC addresses are stored. By default, each learned MAC address is aged out after 300 seconds. The value ranges from 0 to 1,500 minutes. To disable MAC address aging, set the value to zero.

If you want to see the first MAC address that exceeded the learning limit for an interface or VLAN, you can enable the learning-limit violation log for a managed FortiSwitch unit. Only one violation is recorded per interface or VLAN.

By default, logging is disabled. The most recent violation that occurred on each interface or VLAN is recorded in the system log. After that, no more violations are logged until the log is reset for the triggered interface or VLAN. Only the most recent 128 violations are displayed in the console.

Use the following commands to control the learning-limit violation log and to control how long learned MAC addresses are save:

```
config switch-controller global
    set mac-violation-timer <0-1500>
    set log-mac-limit-violations {enable | disable}
end
```

To view the content of the learning-limit violation log for a managed FortiSwitch unit, use one of the following commands:

- `diagnose switch-controller dump mac-limit-violations all <FortiSwitch_serial_number>`
- `diagnose switch-controller dump mac-limit-violations interface <FortiSwitch_serial_number> <port_name>`
- `diagnose switch-controller dump mac-limit-violations vlan <FortiSwitch_serial_number> <VLAN_ID>`

To reset the learning-limit violation log for a managed FortiSwitch unit, use one of the following commands:

- `execute switch-controller mac-limit-violation reset all <FortiSwitch_serial_number>`
- `execute switch-controller mac-limit-violation reset vlan <FortiSwitch_serial_number> <VLAN_ID>`
- `execute switch-controller mac-limit-violation reset interface <FortiSwitch_serial_number> <port_name>`

Testing 802.1x authentication with monitor mode (480807)

Use the monitor mode to test your system configuration for 802.1x authentication. You can use monitor mode to test port-based authentication, MAC-based authentication, EAP pass-through mode, and MAC authentication bypass. Monitor mode is disabled by default. After you enable monitor mode, the network traffic will continue to flow, even if the users fail authentication.

To enable or disable monitor mode, use the following commands:

```
config switch-controller security-policy 802-1X
    edit "<policy_name>"
        set open-auth {enable | disable}
    next
end
```

Checking FortiSwitch connections (453491)

A new CLI command provides detailed diagnostic information on the managed FortiSwitch connections:

```
execute switch-controller diagnose-connection <FortiSwitch_serial_number>
```


If the FortiSwitch serial number is omitted, only the FortiLink configuration is checked.

CLI changes for quarantining MAC addresses (442831)

The CLI commands to create a permanent quarantine of specific MAC addresses have changed to allow multiple MAC addresses to be quarantined in the same session.

```
config user quarantine
  set quarantine enable
  config targets
    edit <quarantine_entry_name>
      set description <string>
      config macs
        edit <MAC_address_1>
        next
        edit <MAC_address_2>
        next
        edit <MAC_address_3>
        next
      end
    end
  end
```

Option	Description
quarantine_entry_name	A name for this quarantine entry.
string	Optional. A description of the MAC addresses being quarantined.
MAC_address_1, MAC_address_2, MAC_address_3	A layer-2 MAC address in the following format: 12:34:56:aa:bb:cc

For example:

```
config user quarantine
  set quarantine enable
  config targets
    edit quarantine1
      config macs
        set description "infected by virus"
        edit 00:00:00:aa:bb:cc
        next
        edit 00:11:22:33:44:55
        next
        edit 00:01:02:03:04:05
        next
      end
    end
  end
```

CLI changes for releasing MAC addresses from quarantine (442831)

To release MAC addresses from quarantine, you can delete a single MAC address or delete a quarantine entry, which will delete all of the MAC addresses listed in the entry.

To delete a single quarantined MAC address:

```
config user quarantine
  config targets
    edit <quarantine_entry_name>
      config macs
        delete <MAC_address_1>
      end
    end
  end
end
```

To delete all MAC addresses in a quarantine entry:

```
config user quarantine
  config targets
    delete <quarantine_entry_name>
  end
end
```

FortiOS 6.0

These features first appeared in FortiOS 6.0.

Limiting the number of learned MAC addresses on a FortiSwitch interface (445087)

You can limit the number of MAC addresses learned on a FortiSwitch interface (port or VLAN). The limit ranges from 1 to 128. If the limit is set to the default value zero, there is no learning limit.

NOTE: Static MAC addresses are not counted in the limit. The limit refers only to learned MAC addresses.

Use the following CLI commands to limit MAC address learning on a VLAN:

```
config switch vlan
  edit <integer>
    set switch-controller-learning-limit <limit>
  end
end
```

For example:

```
config switch vlan
  edit 100
    set switch-controller-learning-limit 20
  end
end
```

Use the following CLI commands to limit MAC address learning on a port:

```
config switch-controller managed-switch
  edit <FortiSwitch_Serial_Number>
    config ports
      edit <port>
        set learning-limit <limit>
      end
    end
  end
end
```

```

        next
    end
end
end

```

For example:

```

config switch-controller managed-switch
  edit S524DF4K15000024
    config ports
      edit port3
        set learning-limit 50
      next
    end
  end
end

```

Sharing FortiSwitch ports between VDOMs (391878)

Virtual domains (VDOMs) are a method of dividing a FortiGate unit into two or more virtual units that function as multiple independent units. VDOMs provide separate security domains that allow separate zones, user authentication, security policies, routing, and VPN configurations.

FortiSwitch ports can now be shared between VDOMs.

NOTE: You cannot use the quarantine feature while sharing FortiSwitch ports between VDOMs.

To share FortiSwitch ports between VDOMs:

1. Create one or more VDOMs.
2. Assign VLANs to each VDOM as required.
3. From these VLANs, select one VLAN to be the default VLAN for the ports in the virtual switch:

```

config switch-controller global
  set default-virtual-switch-vlan <VLAN>

```

NOTE: You must execute these commands from the VDOM that the default VLAN belongs to.

When you add a new port to the VDOM, the new port will be automatically assigned to the default VLAN. You can reassign the ports to other VLANs later.

4. Create a virtual port pool (VPP) to contain the ports to be shared:

```

config switch-controller virtual-port-pool
  edit <VPP_name>
    description <string>
  next
end

```

NOTE: You must execute these commands from the VDOM that the default VLAN belongs to.

For example:

```

config switch-controller virtual-port-pool
  edit "pool3"
    description "pool for port3"

```

```

    next
end

```

5. Share a FortiSwitch port from the VDOM that the FortiSwitch belongs to with another VDOM or export the FortiSwitch port to a VPP where it can be used by any VDOM:

```

config switch-controller managed-switch
  edit <switch.id>
    config ports
      edit <port_name>
        set {export-to-pool <VPP_name> | export-to <VDOM_name>}
        set export-tags <string1,string2,string3,...>
      next
    end
  next
end

```

NOTE: You must execute these commands from the VDOM that the default VLAN belongs to.

For example, if you want to export a port to the VPP named `pool3`:

```

config switch-controller managed-switch
  edit "S524DF4K15000024"
    config ports
      edit port3
        set export-to-pool "pool3"
        set export-tags "Pool 3"
      next
    end
  next
end

```

For example, if you want to export a port to the VDOM named `vdom3`:

```

config switch-controller managed-switch
  edit "S524DF4K15000024"
    config ports
      edit port3
        set export-to "vdom3"
        set export-tags "VDOM 3"
      next
    end
  next
end

```

6. Request a port in a VPP:

```

execute switch-controller virtual-port-pool request <FortiSwitch_device_ID> <port_name>

```

NOTE: You must execute this command from the VDOM that is requesting the port.

For example:

```

execute switch-controller virtual-port-pool request S524DF4K15000024h port3

```

7. Return a port to a VPP:

```

execute switch-controller virtual-port-pool return <FortiSwitch_device_ID> <port_name>

```

NOTE: You must execute this command from the VDOM that owns the port.

For example:

```
execute switch-controller virtual-port-pool return S524DF4K15000024h port3
```

You can create your own export tags using the following CLI commands:

```
config switch-controller switch-interface-tag
  edit <tag_name>
end
```

Use the following CLI command to list the contents of a specific VPP:

```
execute switch-controller virtual-port-pool show-by-pool <VPP_name>
```

Use the following CLI command to list all VPPs and their contents:

```
execute switch-controller virtual-port-pool show
```

NOTE: Shared ports do not support the following features:

- LLDP
- 802.1x
- STP
- BPDU guard
- Root guard
- DHCP snooping
- IGMP snooping
- QoS
- Port security
- MCLAG

sFlow support (450507)

sFlow is a method of monitoring the traffic on your network to identify areas on the network that might impact performance and throughput. With sFlow, you can export truncated packets and interface counters. FortiSwitch implements sFlow version 5 and supports trunks and VLANs.

NOTE: Because sFlow is CPU intensive, Fortinet does not recommend high rates of sampling for long periods.

sFlow uses packet sampling to monitor network traffic. The sFlow agent captures packet information at defined intervals and sends them to an sFlow collector for analysis, providing real-time data analysis. To minimize the impact on network throughput, the information sent is only a sampling of the data.

The sFlow collector is a central server running software that analyzes and reports on network traffic. The sampled packets and counter information, referred to as flow samples and counter samples, respectively, are sent as sFlow datagrams to a collector. Upon receiving the datagrams, the sFlow collector provides real-time analysis and graphing to indicate the source of potential traffic issues. sFlow collector software is available from a number of third-party software vendors. You must configure a FortiGate policy to transmit the samples from the FortiSwitch unit to the sFlow collector.

sFlow can monitor network traffic in two ways:

- Flow samples—You specify the percentage of packets (one out of n packets) to randomly sample.
- Counter samples—You specify how often (in seconds) the network device sends interface counters.

Use the following CLI commands to specify the IP address and port for the sFlow collector. By default, the IP address is 0.0.0.0, and the port number is 6343.

```
config switch-controller sflow
  collector-ip <x.x.x.x>
  collector-port <port_number>
end
```

Use the following CLI commands to configure sFlow:

```
config switch-controller managed-switch <FortiSwitch_serial_number>
  config ports
    edit <port_name>
      set sflow-sampler <disabled | enabled>
      set sflow-sample-rate <0-99999>
      set sflow-counter-interval <1-255>
    next
  next
end
```

For example:

```
config switch-controller sflow
  collector-ip 1.2.3.4
  collector-port 10
end

config switch-controller managed-switch S524DF4K15000024
  config ports
    edit port5
      set sflow-sampler enabled
      set sflow-sample-rate 10
      set sflow-counter-interval 60
    next
  next
end
```

Restricting the type of frames allowed through IEEE 802.1Q ports (448505)

You can now specify whether each FortiSwitch port discards tagged 802.1Q frames or untagged 802.1Q frames or allows all frames access to the port. By default, all frames have access to each FortiSwitch port.

Use the following CLI commands:

```
config switch-controller managed-switch <SN>
  config ports
    edit <port_name>
      set discard-mode <none | all-tagged | all-untagged>
    next
  next
end
```

Dynamic ARP inspection (DAI) support (462511)

DAI prevents man-in-the-middle attacks and IP address spoofing by checking that packets from untrusted ports have valid IP-MAC-address binding. DAI allows only valid ARP requests and responses to be forwarded.

To use DAI, you must first enable the DHCP-snooping feature, enable DAI, and then enable DAI for each VLAN. By default, DAI is disabled on all VLANs.

After enabling DHCP snooping with the `set switch-controller-dhcp-snooping enable` command, use the following CLI commands to enable DAI and then enable DAI for a VLAN:

```
config system interface
  edit vsw.test
    set switch-controller-arp-inspection <enable | disable>
  end

config switch-controller managed-switch
  edit <sn>
    config ports
      edit <VLAN_ID>
        arp-inspection-trust <untrusted | trusted>
      next
    end
  next
end
```

Use the following CLI command to check DAI statistics for a FortiSwitch unit:

```
diagnose switch arp-inspection stats <FortiSwitch_Serial_Number>
```

Use the following CLI command to delete DAI statistics for a specific VLAN:

```
diagnose switch arp-inspection stats clear <VLAN_ID> <FortiSwitch_Serial_Number>
```

FortiSwitch port mirroring support (457122)

The FortiSwitch unit can send a copy of any ingress or egress packet on a port to egress on another port of the same FortiSwitch unit. The original traffic is unaffected. This process is known as port mirroring and is typically used for external analysis and capture.

Use the following CLI commands to configure FortiSwitch port mirroring:

```
config switch-controller managed-switch
  edit <FortiSwitch_Serial_Number>
    config mirror
      edit <mirror_name>
        set status <active | inactive>
        set dst <port_name>
        set switching-packet <enable | disable>
        set src-ingress <port_name>
        set src-egress <port_name>
      next
    end
  next
end
```

NOTE: The `set status` and `set dst` commands are mandatory for port mirroring.

For example:

```
config switch-controller managed-switch
edit S524DF4K15000024
config mirror
edit 2
set status active
set dst port1
set switching-packet enable
set src-ingress port2 port3
set src-egress port4 port5
next
end
next
```

Quarantining MAC addresses (459525)

NOTE: Previously, this feature used the `config switch-controller quarantine` CLI command.

To create a permanent quarantine of specific MAC addresses, use the following CLI commands:

```
config user quarantine
set quarantine enable
config targets
edit <MAC_address>
set description <string>
set tags <tag1 tag2 tag3 ...>
next
end
end
```

Option	Description
MAC_address	A layer-2 MAC address in the following format: 12:34:56:aa:bb:cc
string	Optional. A description of the MAC address being quarantined.
tag1 tag2 tag3 ...	Optional. A list of arbitrary strings.

For example:

```
config user quarantine
set quarantine enable
config targets
edit 00:00:00:aa:bb:cc
set description "infected by virus"
set tags "quarantined"
next
end
end
```

You can add MAC addresses to be quarantined even when the quarantine feature is disabled. The MAC addresses are only quarantined when the quarantine feature is enabled.

Banning IP addresses (459525)

To temporarily ban an IP address, use the following CLI command:

```
diagnose user ban add src4 <IPv4_address>
```


Previously, this feature used the `diagnose user quarantine` CLI command.

Synchronizing the FortiGate unit with the managed FortiSwitch units (454664)

You can now synchronize the FortiGate unit with the managed FortiSwitch units to check for synchronization errors on each managed FortiSwitch unit.

Use the following command to synchronize the full configuration of a FortiGate unit with the managed FortiSwitch unit:

```
execute switch-controller trigger-config-sync <FortiSwitch_serial_number>
```

Use one of the following commands to display the synchronization state of a FortiGate unit with a specific managed FortiSwitch unit:

```
execute switch-controller get-sync-status switch-id <FortiSwitch_serial_number>
execute switch-controller get-sync-status name <FortiSwitch_name>
```

Use the following command to display the synchronization state of a FortiGate unit with a group of managed FortiSwitch units:

```
execute switch-controller get-sync-status group <FortiSwitch_group_name>
```

Use the following command to check the synchronization state of all managed FortiSwitch units in the current VDOM:

```
execute switch-controller get-sync-status all
```

For example:

```
FG100D3G14813513 (root) # execute switch-controller get-sync-status all
Managed-devices in current vdom root:

STACK-NAME: FortiSwitch-Stack-port5
SWITCH (NAME)                STATUS CONFIG      MAC-SYNC      UPGRADE
FS1D243Z14000173            Up      Idle           Idle          Idle
S124DP3X16006228 (Desktop-Switch) Up      Idle           Idle          Idle
```

Enabling the use of HTTPS to download firmware to managed FortiSwitch units (454664)

Use the following CLI commands to enable the use of HTTPS to download firmware to managed FortiSwitch units:

```
config switch-controller global
  set https-image-push enable
end
```

RADIUS accounting support (451023)

The FortiSwitch unit uses 802.1x-authenticated ports to send five types of RADIUS accounting messages to the RADIUS accounting server to support FortiGate RADIUS single sign-on:

- START—The FortiSwitch has been successfully authenticated, and the session has started.
- STOP—The FortiSwitch session has ended.
- INTERIM—Periodic messages sent based on the value set using the `set acct-interim-interval` command.
- ON—FortiSwitch will send this message when the switch is turned on.
- OFF—FortiSwitch will send this message when the switch is shut down.

Use the following commands to set up RADIUS accounting so that FortiOS can send accounting messages to managed FortiSwitch units:

```
config user radius
  edit <RADIUS_server_name>
    set acct-interim-interval <seconds>
  config accounting-server
    edit <entry_ID>
      set status {enable | disable}
      set server <server_IP_address>
      set secret <secret_key>
      set port <port_number>
    next
  end
next
end
```

FortiLink mode supported over a layer-3 network (457103)

This feature allows FortiSwitch islands (FSIs) to operate in FortiLink mode over a layer-3 network, even though they are not directly connected to the switch-controller FortiGate unit. FSIs contain one or more FortiSwitch units.

The following limitations apply to FSIs operating in FortiLink mode over a layer-3 network:

- All FortiSwitch units using this feature must be included in the FortiGate preconfigured switch table.
- No layer-2 data path component, such as VLANs, can span across layer 3 between the FortiGate unit and the FortiSwitch unit.
- All FortiSwitch units within an FSI must be connected to the same FortiGate unit.
- The FortiSwitch unit needs a functioning layer-3 routing configuration to reach the FortiGate unit or any feature-configured destination, such as syslog or 802.1x.
- Do not connect a layer-2 FortiGate unit and a layer-3 FortiGate unit to the same FortiSwitch unit.
- If the FortiSwitch management port is used for a layer-3 connection to the FortiGate unit, the FSI can contain only one FortiSwitch unit. All switch ports must remain in standalone mode.
- Do not connect a FortiSwitch unit to a layer-3 network and a layer-2 network on the same segment.
- If the network has a wide geographic distribution, some features, such as software downloads, might operate slowly.

To configure a FortiSwitch unit to operate in a layer-3 network:

1. Reset the FortiSwitch to factory default settings with the `execute factoryreset` command.
2. Manually set the FortiSwitch unit to FortiLink mode:

```
config system global
    set switch-mgmt-mode fortilink
end
```

3. Configure the discovery setting for the FortiSwitch unit. You can either use DHCP discovery or static discovery.

To use DHCP discovery:

```
config switch-controller global
    ac-discovery dhcp
        dhcp-option-code <integer>
    end
end
```

To use static discovery:

```
config switch-controller global
    ac-discovery static
        config ac-list
            id <integer>
                set ipv4-address <IPv4_address>
            next
        end
    end
end
```

4. Configure at least one port of the FortiSwitch unit as an uplink port. When the FortiSwitch is in FortiLink mode, VLAN 4094 is configured on an internal port, which can provide a path to the layer-3 network with the following commands:

```
config switch interface
    edit <port_number>
        set fortilink-l3-mode enable
    end
end
```

NOTE: The NTP server must be configured on the FortiSwitch unit either manually or provided by DHCP. The NTP server must be reachable from the FortiSwitch unit.

Limiting the number of parallel process for FortiSwitch configuration (457103)

Use the following CLI commands to reduce the number of parallel process that the switch controller uses for configuring FortiSwitch units:

```
config global
    config switch-controller system
        set parallel-process-override enable
        set parallel-process <1-300>
```

```
end
end
```

CLI changes for FortiLink mode (447349, 473773)

There are changes to the `execute switch-controller get-physical-connection`, `execute switch-controller get-conn-status`, and `diagnose switch-controller dump network-upgrade status` CLI commands.

- The `execute switch-controller get-physical-connection` CLI command has new parameters:

Use the `execute switch-controller get-physical-connection standard` command to get the FortiSwitch stack connectivity graph in the standard output format.

Use the `execute switch-controller get-physical-connection dot` command to get the FortiSwitch stack connectivity graph in a .dot (Graphviz) output format.

- The `execute switch-controller get-conn-status` CLI command output now includes virtual FortiSwitch units. Virtual FortiSwitch units are indicated by an asterisk (*) after the switch identifier. For example:

```
execute switch-controller get-conn-status

STACK-NAME: FortiSwitch-Stack-port2
SWITCH-ID      VERSION  STATUS      ADDRESS      JOIN-TIME      NAME
S108DV2EJZDAC42F v3.6.0   Authorized/Up 169.254.2.4   Thu Feb 8 17:07:35 2018 -
S108DV4FQON40Q07 v3.6.0   Authorized/Up 169.254.2.5   Thu Feb 8 17:08:37 2018 -
S108DVBWVLH4QGEB v3.6.0   Authorized/Up 169.254.2.6   Thu Feb 8 17:09:13 2018 -
S108DVCY19SA0CD8 v3.6.0   Authorized/Up 169.254.2.2   Thu Feb 8 17:04:41 2018 -
S108DVD98KMQGC44* v3.6.0   Authorized/Up 169.254.2.7   Thu Feb 8 17:10:50 2018 -
S108DVGGBJLQQO48* v3.6.0   Authorized/Up 169.254.2.3   Thu Feb 8 17:06:57 2018 -
S108DVKM5T2QEA92 v3.6.0   Authorized/Up 169.254.2.8   Thu Feb 8 17:11:00 2018 -
S108DVZX3VTA0045 v3.6.0   Authorized/Up 169.254.2.9   Thu Feb 8 17:11:00 2018 -

Managed-Switches: 8      UP: 8      DOWN: 0
```

- The `diagnose switch-controller dump network-upgrade status` CLI command output now includes the location of the image that is loaded when the FortiSwitch unit is restarted. If the Next boot column is blank, the FortiSwitch unit uses the same location each it is restarted. The status column shows the percentage downloaded, the percentage erased in flash memory, and the percentage written to flash memory.

For example:

```
diagnose switch-controller dump network-upgrade status

Running                                     Status      Next boot
-----
VDOM : root
S108DVCY19SA0CD8 S108DV-v3.6.0-build4277,171207 (Interim) (0/0/0) S108DV-v3.7.0-
build4277,171207 (Interim)
S108DV2EJZDAC42F S108DV-v3.6.0-build4277,171207 (Interim) (0/0/0)
```

Upgrade the firmware on multiple FortiSwitch units at the same time using the GUI (462553)

To upgrade the firmware on multiple FortiSwitch units at the same time:

- Go to *WiFi & Switch Controller > Managed FortiSwitch*.
 - Select the faceplates of the FortiSwitch units that you want to upgrade.
 - Click *Upgrade*.
- The *Upgrade FortiSwitches* page opens.

4. Select *FortiGuard* or select *Upload* and then select the firmware file to upload.
You can select only one firmware image to use to upgrade the selected FortiSwitch units. If the FortiSwitch unit already has the latest firmware image, it will not be upgraded.
5. Select *Upgrade*.

Network-assisted device detection (377467)

Network-assisted device detection allows the FortiGate unit to use the information about connected devices detected by the managed FortiSwitch unit.

To enable network-assisted device detection on a VDOM:

```
config switch-controller network-monitor-settings
    set network-monitoring enable
end
```

Connecting FortiLink ports

This section contains information about the FortiSwitch and FortiGate ports that you connect to establish a FortiLink connection.

In FortiSwitchOS 3.3.0 and later releases, you can use any of the switch ports for FortiLink. Some or all of the switch ports (depending on the model) support auto-discovery of the FortiLink ports.

You can choose to connect a single FortiLink port or multiple FortiLink ports as a logical interface (link-aggregation group, hardware switch, or software switch).

1. Enable the switch controller on the FortiGate unit

Before connecting the FortiSwitch and FortiGate units, ensure that the switch controller feature is enabled on the FortiGate unit with the FortiGate web-based manager or CLI to enable the switch controller. Depending on the FortiGate model and software release, this feature might be enabled by default.

Using the FortiGate GUI

1. Go to *System > Feature Visibility*.
2. Turn on the *Switch Controller* feature, which is in the *Basic Features* list.
3. Select *Apply*.

The menu option *WiFi & Switch Controller* now appears.

Using the FortiGate CLI

Use the following commands to enable the switch controller:

```
config system global
    set switch-controller enable
end
```

2. Connect the FortiSwitch unit and FortiGate unit

FortiSwitchOS 3.3.0 and later provides flexibility for FortiLink:

- Use any switch port for FortiLink
- Provides auto-discovery of the FortiLink ports on the FortiSwitch
- Choice of a single FortiLink port or multiple FortiLink ports in a link-aggregation group (LAG)

Auto-discovery of the FortiSwitch ports

In FortiSwitchOS 3.3.0 and later releases, D-series FortiSwitch models support FortiLink auto-discovery, on automatic detection of the port connected to the FortiGate unit.

You can use any of the switch ports for FortiLink. Before connecting the switch to the FortiGate unit, use the following FortiSwitch CLI commands to configure a port for FortiLink auto-discovery:

```
config switch interface
    edit <port>
        set auto-discovery-fortilink enable
    end
end
```

end

By default, each FortiSwitch model provides a set of ports that are enabled for FortiLink auto-discovery. If you connect the FortiLink using one of these ports, no switch configuration is required.

In FortiSwitchOS 3.4.0 and later releases, the last four ports are the default auto-discovery FortiLink ports. You can also run the `show switch interface` command on the FortiSwitch unit to see the ports that have auto-discovery enabled.

The following table lists the default auto-discovery ports for each switch model.

NOTE: Any port can be used for FortiLink if it is manually configured.

FortiSwitch Model	Default Auto-FortiLink ports
FS-108D	ports 9 and 10
FS-108D-POE	ports 9 and 10
FSR-112D	ports 9, 10, 11 and 12
FSR-112D-POE	ports 5, 6, 7, 8, 9, 10, 11, and 12
FS-124D, FS-124D-POE	ports 23, 24, 25, and 26
FS-224D-POE	ports 21, 22, 23, and 24
FS-224D-FPOE	ports 21, 22, 23, 24, 25, 26, 27, and 28
FS-248D, FS-248D-FPOE, FS-448D, FS-448D-FPOE, FS-448D-POE	ports 45, 46, 47, 48, 49, 50, 51, and 52
FS-248D-POE	ports 47, 48, 49, and 50
FS-424D, FS-424D-POE, FS-424D-FPOE	ports 23, 24, 25, and 26
FS-524D, FS-524D-FPOE	ports 21, 22, 23, 24, 25, 26, 27, 28, 29, and 30
FS-548D, FS-548D-FPOE	ports 45, 46, 47, 48, 49, 50, 51, 52, 53, and 54
FS-1024D, FS-1048D, FS-3032D	all ports

Choosing the FortiGate ports

The FortiGate unit manages all of the switches through one active FortiLink. The FortiLink can consist of one port or multiple ports (for a LAG).

As a general rule, FortiLink is supported on all ports that are not listed as HA ports.

FortiLink configuration using the FortiGate GUI

This section describes how to configure a FortiLink between a FortiSwitch unit and a FortiGate unit.

You can configure FortiLink using the FortiGate GUI or CLI. Fortinet recommends using the GUI because the CLI procedures are more complex (and therefore more prone to error).

If you use one of the auto-discovery FortiSwitch ports, you can establish the FortiLink connection (single port or LAG) with no configuration steps on the FortiSwitch and with a few simple configuration steps on the FortiGate unit.

Summary of the procedure

1. On the FortiGate unit, configure the FortLink port or create a logical FortLink interface.
2. Authorize the managed FortiSwitch unit.

Configure FortiLink as a single link

To configure the FortiLink port on the FortiGate unit:

1. Go to *Network > Interfaces*.
2. (Optional) If the FortiLink physical port is currently included in the internal interface, edit it and remove the desired port from the Physical Interface Members.
3. Edit the FortiLink port.
4. Set *Addressing mode* to *Dedicated to FortiSwitch*.
5. Configure the *IP/Network Mask* for your network.
6. Optionally select *Automatically authorize devices* or disable to manually authorize the FortiSwitch.
7. Select *OK*.

Configure FortiLink as a logical interface

You can configure the FortiLink as a logical interface: link-aggregation group (LAG), hardware switch, or software switch).

LAG is supported on all FortiSwitch models and on FortiGate models FGT-100D and above. Hardware switch is supported on some FortiGate models.

Connect any of the FortiLink-capable ports on the FortiGate unit to the FortiSwitch unit. Ensure that you configure auto-discovery on the FortiSwitch ports (unless it is so by default).

1. Go to *Network > Interfaces*.
2. (Optional) If the FortiLink physical ports are currently included in the internal interface, edit the internal interface, and remove the desired ports from the Physical Interface Members.
3. Select *Create New > Interface*.
4. Enter a name for the interface (11 characters maximum).
5. Set the *Type* to *802.3ad Aggregate*, *Hardware Switch*, or *Software Switch*.
6. Select the FortiGate ports for the logical interface.
7. Set *Addressing mode* to *Dedicated to FortiSwitch*.
8. Configure the *IP/Network Mask* for your network.

9. Optionally select *Automatically authorize devices* or disable to manually authorize the FortiSwitch.
10. Select *OK*.

FortiLink split interface

You can use the FortiLink split interface to connect the FortiLink aggregate interface from one FortiGate unit to two FortiSwitch units. When the FortiLink split interface is enabled, only one link remains active.

The aggregate interface for this configuration must contain exactly two physical ports (one for each FortiSwitch unit).

You must enable the split interface on the FortiLink aggregate interface using the FortiGate CLI:

```
config system interface
    edit <name of the FortiLink interface>
        set fortilink-split-interface enable
    end
```

Authorizing the FortiSwitch unit

If you configured the FortiLink interface to manually authorize the FortiSwitch unit as a managed switch, perform the following steps:

1. Go to *WiFi & Switch Controller > Managed FortiSwitch*.
2. Optionally, click on the FortiSwitch faceplate and click *Authorize*. This step is required only if you disabled the automatic authorization field of the interface.

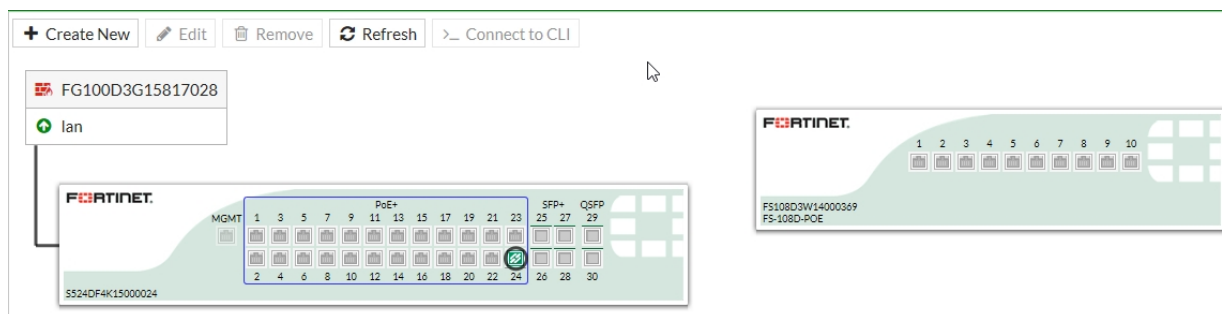
Adding preauthorized FortiSwitch units

After you preauthorize a FortiSwitch unit, you can assign the FortiSwitch ports to a VLAN.

To preauthorize a FortiSwitch:

1. Go to *WiFi & Switch Controller > Managed FortiSwitch*.
2. Click *Create New*.
3. In the New Managed FortiSwitch page, enter the serial number, model name, and description of the FortiSwitch.
4. Move the *Authorized* slider to the right.
5. Click *OK*.

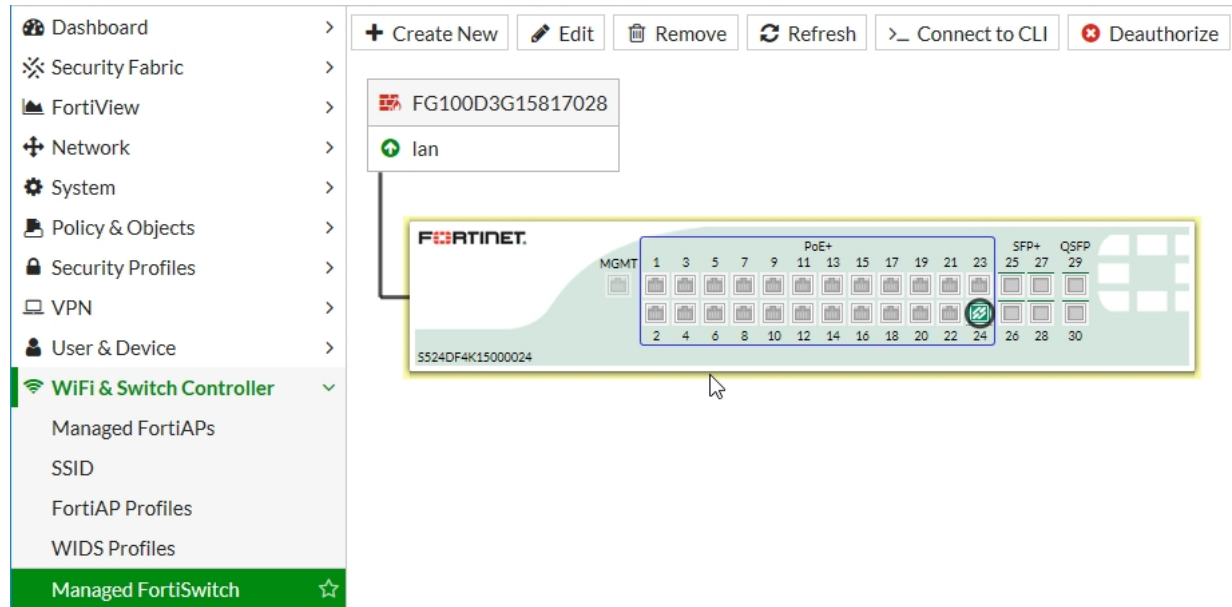
The Managed FortiSwitch page shows a FortiSwitch faceplate for the preauthorized switch.



Managed FortiSwitch display

Go to *WiFi & Switch Controller > Managed FortiSwitch* to see all of the switches being managed by your FortiGate.

When the FortiLink is established successfully, the status is green (next to the FortiGate interface name and on the FortiSwitch faceplate), and the link between the ports is a solid line.



If the link has gone down for some reason, the line will be dashed, and a broken link icon will appear. You can still edit the FortiSwitch unit though and find more information about the status of the switch. The link to the FortiSwitch unit might be down for a number of reasons; for example, a problem with the cable linking the two devices, firmware versions being out of synch, and so on. You need to make sure the firmware running on the FortiSwitch unit is compatible with the firmware running on the FortiGate unit.

From the Managed FortiSwitch page, you can edit any of the managed FortiSwitch units, remove a FortiSwitch unit from the configuration, refresh the display, connect to the CLI of a FortiSwitch unit, or deauthorize a FortiSwitch unit.

Edit a managed FortiSwitch unit

To edit a managed FortiSwitch unit:

1. Go to *Wifi & Switch Controller > Managed FortiSwitch*.
2. Click on the FortiSwitch to and click *Edit*, right-click on a FortiSwitch unit and select *Edit*, or double-click on a FortiSwitch unit.

From the *Edit Managed FortiSwitch* form, you can:

- Change the *Name* and *Description* of the FortiSwitch unit.
- View the *Status* of the FortiSwitch unit.
- *Restart* the FortiSwitch.

- *Authorize* or deauthorize the FortiSwitch.
- *Update* the firmware running on the switch.

Network interface display

On the *Network > Interfaces* page, you can see the FortiGate interface connected to the FortiSwitch unit. The GUI indicates *Dedicated to FortiSwitch* in the IP/Netmask field.

+

Create New

✎

Edit

🗑

Delete

By Type

By Role

Alphabetically

▼ Status	▼ Name	▼ Members	▼ IP/Netmask	▼ Type	▼ Access	▼ Ref.
Physical (4)						
🟢	port1		172.20.121.31 255.255.255.0	Physical Interface	PING HTTPS	2
🟢	port2		1.1.1.1 255.255.255.0	Physical Interface		2
🟢	port3 (1 Connected FortiSwitch(s))		Dedicated to FortiSwitch	Physical Interface	PING CAPWAP	3
	vsw.port3		0.0.0.0 0.0.0.0	VLAN		10

Add link aggregation groups (Trunks)

To create a link aggregation group for FortiSwitch user ports:

1. Go to *WiFi & Switch Controller > FortiSwitch Ports*.
2. Click *Create New > Trunk*.
3. In the New Trunk Group page, enter a *Name* for the trunk group.
4. Select two or more physical ports to add to the trunk group.
5. Select the *Mode*: Static, Passive LACP, or Active LACP.
6. Click OK.

New Trunk Group

Name:

Members: 🔴 port1 ✕ 🔴 port2 ✕ 🔴 port3 ✕ +

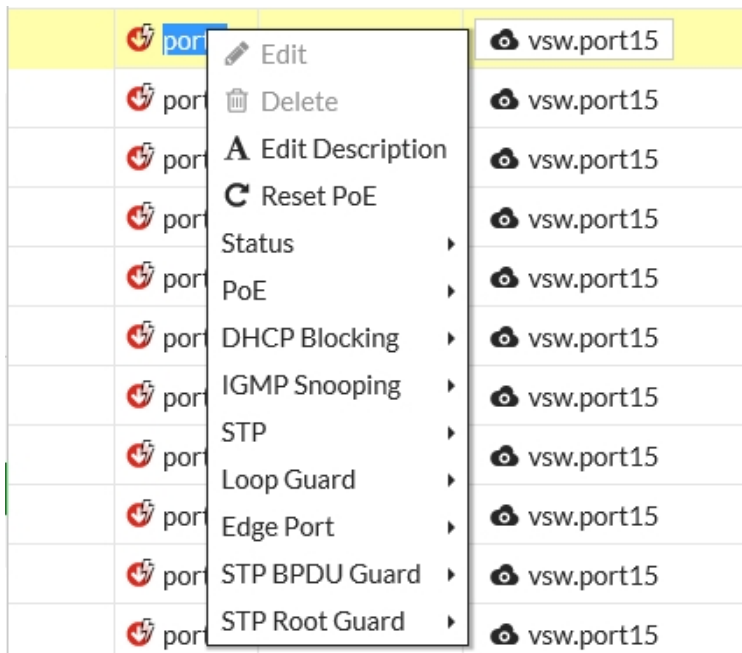
Mode: Static Passive LACP Active LACP

Configure DHCP blocking, IGMP snooping, STP, and loop guard on managed FortiSwitch ports

Go to *WiFi & Switch Controller > FortiSwitch Ports*. Right-click any port and then enable or disable the following features:

- **DHCP blocking**—The DHCP blocking feature monitors the DHCP traffic from untrusted sources (for example, typically host ports and unknown DHCP servers) that might initiate traffic attacks or other hostile actions. To prevent this, DHCP blocking filters messages on untrusted ports.
- **IGMP snooping**—IGMP snooping allows the FortiSwitch to passively listen to the Internet Group Management Protocol (IGMP) network traffic between hosts and routers. The switch uses this information to determine which ports are interested in receiving each multicast feed. FortiSwitch can reduce unnecessary multicast traffic on the LAN by pruning multicast traffic from links that do not contain a multicast listener.
- **Spanning Tree Protocol (STP)**—STP is a link-management protocol that ensures a loop-free layer-2 network topology.
- **Loop guard**—A loop in a layer-2 network results in broadcast storms that have far-reaching and unwanted effects. Fortinet loop guard helps to prevent loops. When loop guard is enabled on a switch port, the port monitors its subtending network for any downstream loops. The loop guard feature is designed to work in concert with STP rather than as a replacement for STP.
- **STP root guard**—Root guard protects the interface on which it is enabled from becoming the path to root. When enabled on an interface, superior BPDUs received on that interface are ignored or dropped. Without using root guard, any switch that participates in STP maintains the ability to reroute the path to root. Rerouting might cause your network to transmit large amounts of traffic across suboptimal links or allow a malicious or misconfigured device to pose a security risk by passing core traffic through an insecure device for packet capture or inspection. By enabling root guard on multiple interfaces, you can create a perimeter around your existing paths to root to enforce the specified network topology.
- **STP BPDU guard**—Similar to root guard, BPDU guard protects the designed network topology. When BPDU guard is enabled on STP edge ports, any BPDUs received cause the ports to go down for a specified number of minutes. The BPDUs are not forwarded, and the network edge is enforced.

STP is enabled on all ports by default. Loop guard is disabled by default on all ports.



FortiLink configuration using the FortiGate CLI

This section describes how to configure FortiLink using the FortiGate CLI. Fortinet recommends using the FortiGate GUI because the CLI procedures are more complex (and therefore more prone to error).

If you use one of the auto-discovery FortiSwitch ports, you can establish the FortiLink connection (single port or LAG) with no configuration steps on the FortiSwitch and with a few simple configuration steps on the FortiGate unit.

You can also configure FortiLink mode over a layer-3 network.

Summary of the procedure

1. Configure FortiLink on a physical port or configure FortiLink on a logical interface.
2. Configure NTP.
3. Authorize the managed FortiSwitch unit.
4. Configure DHCP.

Configure FortiLink on a physical port

Configure FortiLink on any physical port on the FortiGate unit and authorize the FortiSwitch unit as a managed switch.

In the following steps, port 1 is configured as the FortiLink port.

1. If required, remove port 1 from the `lan` interface:

```
config system virtual-switch
  edit lan
    config port
      delete port1
    end
  end
end
```

2. Configure port 1 as the FortiLink interface:

```
config system interface
  edit port1
    set auto-auth-extension-device enable
    set fortilink enable
  end
end
```

3. Configure an NTP server on port 1:

```
config system ntp
  set server-mode enable
  set interface port1
end
```

4. Authorize the FortiSwitch unit as a managed switch.

```
config switch-controller managed-switch
  edit FS224D3W14000370
```

```

        set fsw-wan1-admin enable
    end
end

```

NOTE: The FortiSwitch unit will reboot when you issue the `set fsw-wan1-admin enable` command.

Configure FortiLink on a logical interface

You can configure FortiLink on a logical interface: link-aggregation group (LAG), hardware switch, or software switch).

NOTE: LAG is supported on all FortiSwitch models and on FortiGate models FGT-100D and above. Hardware switch is supported on some FortiGate models.

Connect any of the FortiLink-capable ports on the FortiGate to the FortiSwitch. Ensure that you configure auto-discovery on the FortiSwitch ports (unless it is auto-discovery by default).

In the following procedure, port 4 and port 5 are configured as a FortiLink LAG.

1. If required, remove the FortiLink ports from the **lan** interface:

```

config system virtual-switch
    edit lan
        config port
            delete port4
            delete port5
        end
    end
end

```

2. Create a trunk with the two ports that you connected to the switch:

```

config system interface
    edit flink1 (enter a name, 11 characters maximum)
        set ip 169.254.3.1 255.255.255.0
        set allowaccess ping capwap https
        set vlanforward enable
        set type aggregate
        set member port4 port5
        set lacp-mode static
        set fortilink enable
        (optional) set fortilink-split-interface enable
    next
end

```

NOTE: If the members of the aggregate interface connect to more than one FortiSwitch, you must enable `fortilink-split-interface`.

3. Authorize the FortiSwitch unit as a managed switch.

```

config switch-controller managed-switch
    edit FS224D3W14000370
        set fsw-wan1-admin enable
    end
end

```

NOTE: FortiSwitch will reboot when you issue the `set fsw-wan1-admin enable` command.

Enable multiple FortiLink interfaces

NOTE: Only the first FortiLink interface has GUI support.

Use the following command to enable or disable multiple FortiLink interfaces.

```
config switch-controller global
    set allow-multiple-interfaces {enable | disable}
end
```

FortiLink mode over a layer-3 network

This feature allows FortiSwitch islands (FSIs) to operate in FortiLink mode over a layer-3 network, even though they are not directly connected to the switch-controller FortiGate unit. FSIs contain one or more FortiSwitch units.

The following limitations apply to FSIs operating in FortiLink mode over a layer-3 network:

- All FortiSwitch units using this feature must be included in the FortiGate preconfigured switch table.
- No layer-2 data path component, such as VLANs, can span across layer 3 between the FortiGate unit and the FortiSwitch unit.
- All FortiSwitch units within an FSI must be connected to the same FortiGate unit.
- The FortiSwitch unit needs a functioning layer-3 routing configuration to reach the FortiGate unit or any feature-configured destination, such as syslog or 802.1x.
- Do not connect a layer-2 FortiGate unit and a layer-3 FortiGate unit to the same FortiSwitch unit.
- If the FortiSwitch management port is used for a layer-3 connection to the FortiGate unit, the FSI can contain only one FortiSwitch unit. All switch ports must remain in standalone mode.
- Do not connect a FortiSwitch unit to a layer-3 network and a layer-2 network on the same segment.
- If the network has a wide geographic distribution, some features, such as software downloads, might operate slowly.

To configure a FortiSwitch unit to operate in a layer-3 network:

1. Reset the FortiSwitch to factory default settings with the `execute factoryreset` command.
2. Manually set the FortiSwitch unit to FortiLink mode:

```
config system global
    set switch-mgmt-mode fortilink
end
```

3. Configure the discovery setting for the FortiSwitch unit. You can either use DHCP discovery or static discovery.

To use DHCP discovery:

```
config switch-controller global
    ac-discovery dhcp
        dhcp-option-code <integer>
    end
end
```

To use static discovery:

```
config switch-controller global
    ac-discovery static
        config ac-list
            id <integer>
                set ipv4-address <IPv4_address>
            next
        end
    end
end
```

4. Configure at least one port of the FortiSwitch unit as an uplink port. When the FortiSwitch is in FortiLink mode, VLAN 4094 is configured on an internal port, which can provide a path to the layer-3 network with the following commands:

```
config switch interface
    edit <port_number>
        set fortilink-l3-mode enable
    end
end
```

NOTE: The NTP server must be configured on the FortiSwitch unit either manually or provided by DHCP. The NTP server must be reachable from the FortiSwitch unit.

Network topologies for managed FortiSwitch units

The FortiGate unit requires only one active FortiLink to manage all of the subtending FortiSwitch units (called *stacking*).

You can configure the FortiLink as a physical interface or as a logical interface (associated with one or more physical interfaces). Depending on the network topology, you can also configure a standby FortiLink.

NOTE: For any of the topologies:

- All of the managed FortiSwitch units will function as one Layer-2 stack where the FortiGate unit manages each FortiSwitch separately.
- The active FortiLink carries data as well as management traffic.

Supported topologies

Fortinet recommends the following topologies for managed FortiSwitch units:

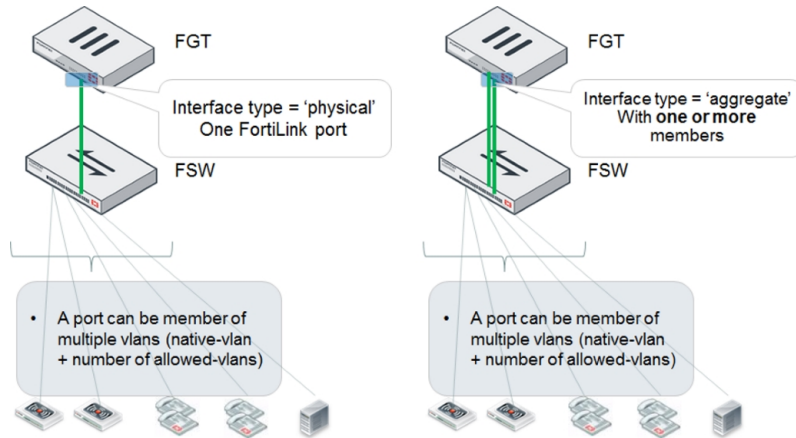
- [Single FortiGate managing a single FortiSwitch unit on page 1931](#)
- [Single FortiGate unit managing a stack of several FortiSwitch units on page 1932](#)
- [HA-mode FortiGate units managing a single FortiSwitch unit on page 1933](#)
- [HA-mode FortiGate units managing a stack of several FortiSwitch units on page 1934](#)
- [HA-mode FortiGate units managing a FortiSwitch two-tier topology on page 1935](#)
- [Single FortiGate unit managing multiple FortiSwitch units \(using a hardware or software switch interface\) on page 1936](#)
- [HA-mode FortiGate units managing two-tier FortiSwitch units with access rings on page 1937](#)
- [Dual-homed servers connected to FortiLink tier-1 FortiSwitch units using an MCLAG on page 1938](#)
- [Standalone FortiGate unit with dual-homed FortiSwitch access on page 1939](#)
- [HA-mode FortiGate units with dual-homed FortiSwitch access on page 1940](#)
- [Multi-tiered MCLAG with HA-mode FortiGate units on page 1941](#)

Single FortiGate managing a single FortiSwitch unit

On the FortiGate unit, the FortiLink interface is configured as physical or aggregate. The 802.3ad aggregate interface type provides a logical grouping of one or more physical interfaces.

NOTE:

- For the aggregate interface, you must disable the split interface on the FortiGate unit.
- When you are using the aggregate interface on the FortiGate unit for the FortiLink interface, the `lacp-mode` of the FortiLink aggregate interface must be set to `static`.



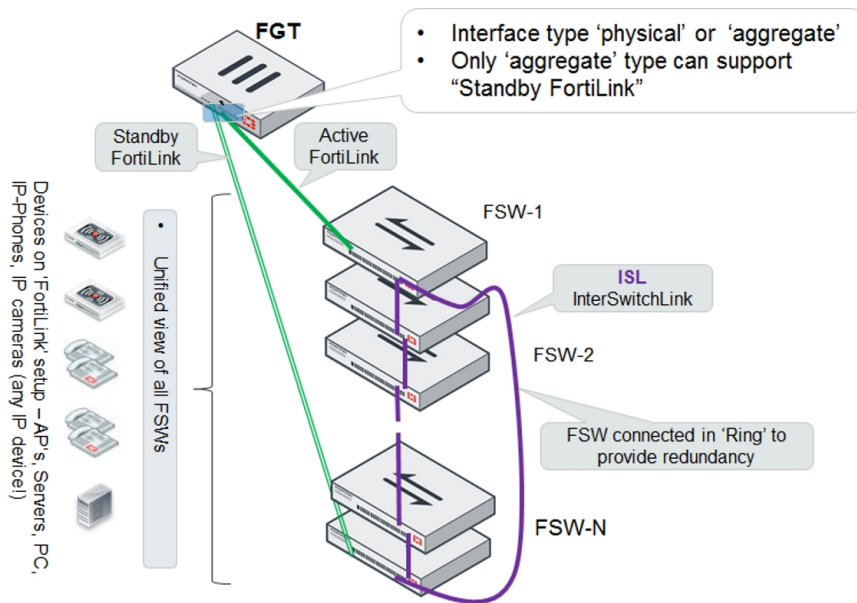
Single FortiGate unit managing a stack of several FortiSwitch units

The FortiGate unit connects directly to one FortiSwitch unit using a physical or aggregate interface. The remaining FortiSwitch units connect in a ring using inter-switch links (that is, ISL).

Optionally, you can connect a standby FortiLink connection to the last FortiSwitch unit. For this configuration, you create a FortiLink Split-Interface (an aggregate interface that contains one active link and one standby link).

NOTE:

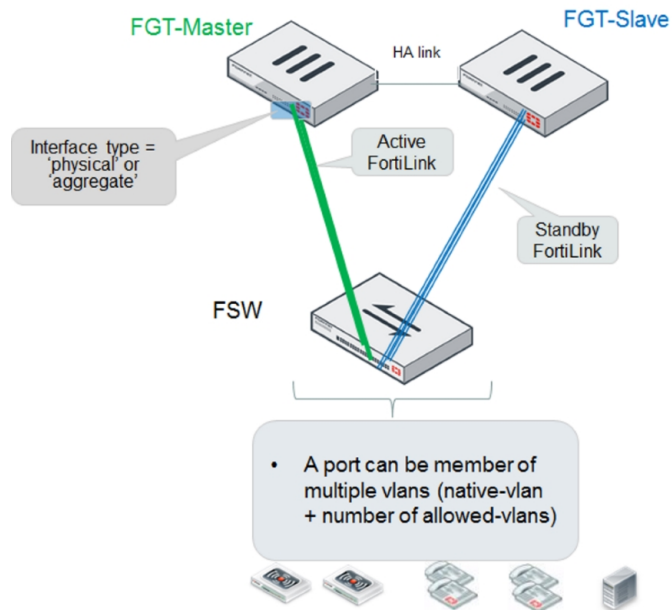
- When you are using the aggregate interface on the FortiGate unit for the FortiLink interface, the `lacp-mode` of the FortiLink aggregate interface must be set to `static`.
- External devices shown in the following topology must be compliant endpoints, such as computers. They cannot be third-party switches or appliances.



HA-mode FortiGate units managing a single FortiSwitch unit

The master and slave FortiGate units both connect a FortiLink to the FortiSwitch unit. The FortiLink port(s) and interface type must match on the two FortiGate units.

NOTE: When using HA-mode FortiGate units to manage FortiSwitch units, the HA mode must be active-passive.




HA-mode FortiGate units managing a stack of several FortiSwitch units

The master and slave FortiGate units both connect a FortiLink to the first FortiSwitch unit and (optionally) to the last FortiSwitch unit. The FortiLink ports and interface type must match on the two FortiGate units.

For the active/standby FortiLink configuration, you create a FortiLink Split-Interface (an aggregate interface that contains one active link and one standby link).

NOTE:

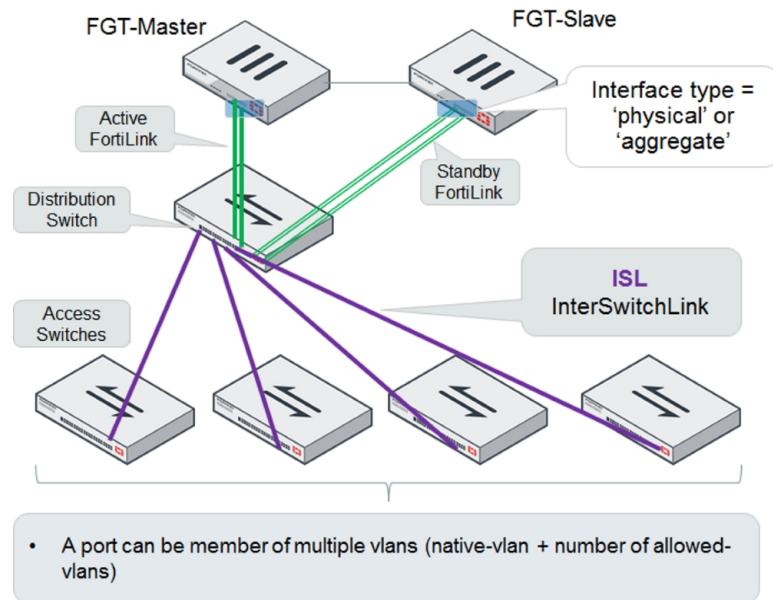
- When you are using the aggregate interface on the FortiGate unit for the FortiLink interface, the `lacp-mode` of the FortiLink aggregate interface must be set to `static`.
- When using HA-mode FortiGate units to manage FortiSwitch units, the HA mode must be active-passive.



HA-mode FortiGate units managing a FortiSwitch two-tier topology

The distribution FortiSwitch unit connects to the master and slave FortiGate units. The FortiLink port(s) and interface type must match on the two FortiGate units.

NOTE: When using HA-mode FortiGate units to manage FortiSwitch units, the HA mode must be active-passive.

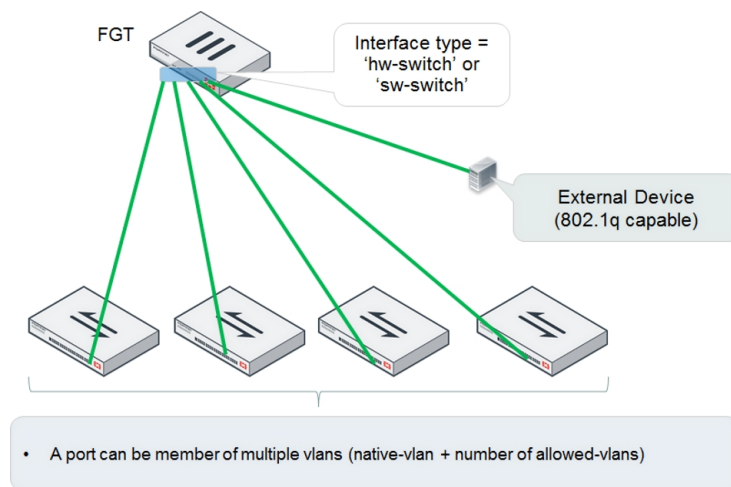


Single FortiGate unit managing multiple FortiSwitch units (using a hardware or software switch interface)

The FortiGate unit connects directly to each FortiSwitch unit. Each of these FortiLink ports is added to the logical hardware-switch or software-switch interface on the FortiGate unit.

Optionally, you can connect other devices to the FortiGate logical interface. These devices, which must support IEEE 802.1q VLAN tagging, will have Layer 2 connectivity with the FortiSwitch ports.

NOTE: Using the hardware or software switch interface in FortiLink mode is not recommended in most cases. It can be used when the traffic on the ports is very light because all traffic across the switches moves through the FortiGate unit.



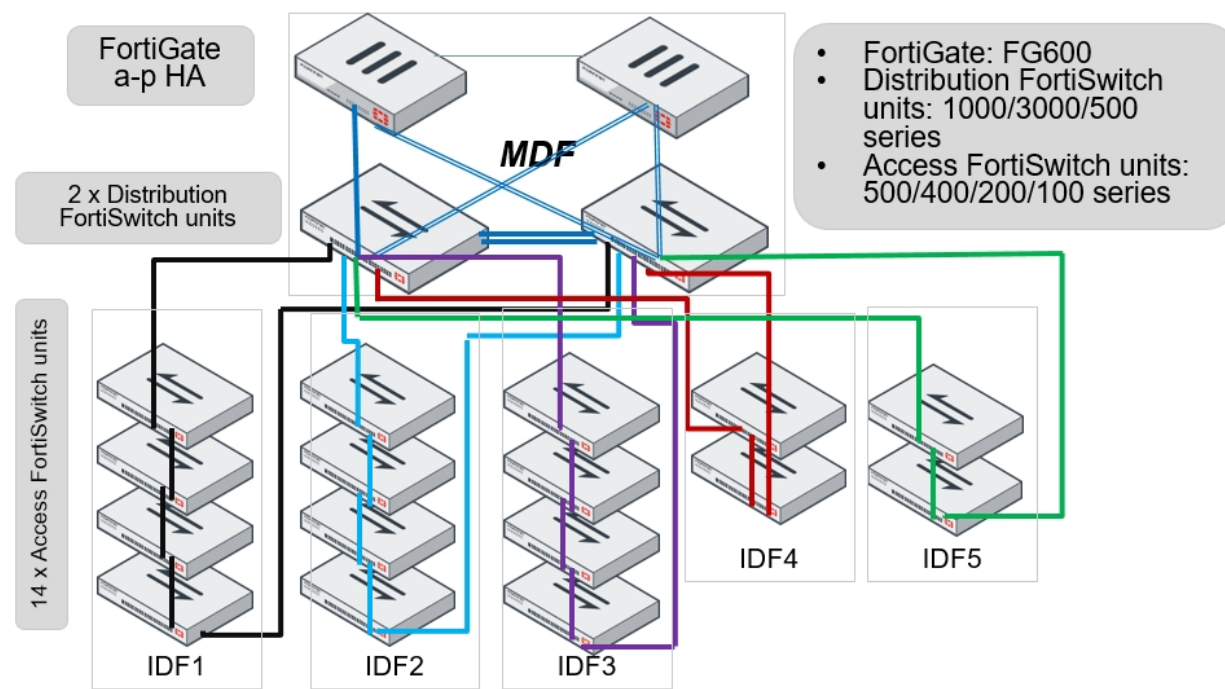
HA-mode FortiGate units managing two-tier FortiSwitch units with access rings

HA-mode FortiGate units connect to redundant distribution FortiSwitch units. Access FortiSwitch units are arranged in a stack in each IDF, connected to both distribution switches.

For the FortiLink connection to each distribution switch, you create a FortiLink split interface (an aggregate interface that contains one active link and one standby link).

NOTE:

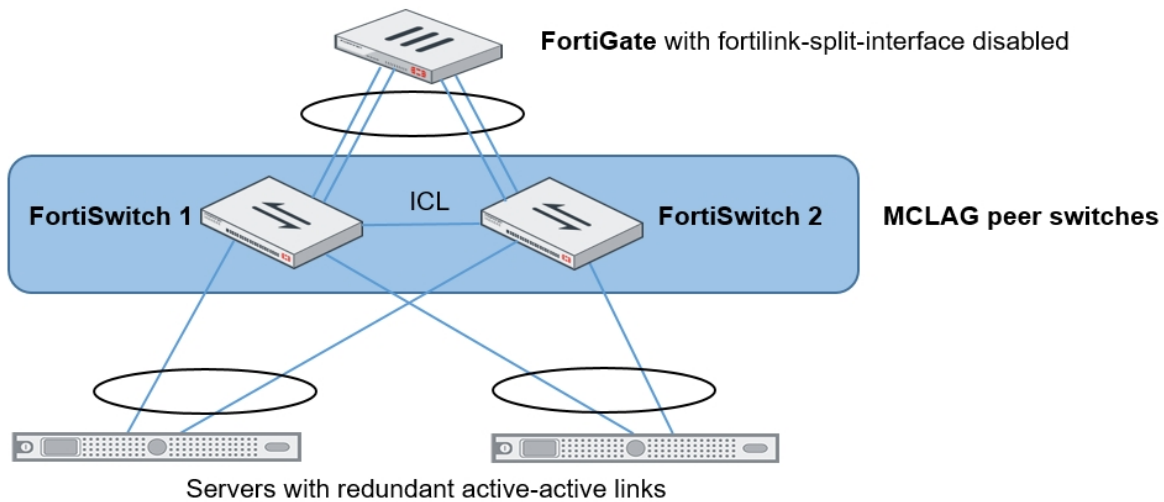
- Before FortiSwitchOS 3.6.4, MCLAG was not supported when access rings were present. Starting with FortiSwitchOS 3.6.4, MCLAG is supported, even with access rings present.
- When using HA-mode FortiGate units to manage FortiSwitch units, the HA mode must be active-passive.
- When you are using the aggregate interface on the FortiGate unit for the FortiLink interface, the `lacp-mode` of the FortiLink aggregate interface must be set to `static`.
- This is only an example topology. Other combinations of FortiGate units and FortiSwitch units can be used to create a similar topology.



Dual-homed servers connected to FortiLink tier-1 FortiSwitch units using an MCLAG

To configure a multichassis LAG, you need to configure FortiSwitch 1 and FortiSwitch 2 as MCLAG peer switches before creating a two-port LAG. Use the `set mclag-icl enable` command to create an inter-chassis link (ICL) on each FortiSwitch unit. Then you set up two MCLAGs towards the servers, each MCLAG using one port from each FortiSwitch unit. You must disable the FortiLink split interface for the FortiGate unit.

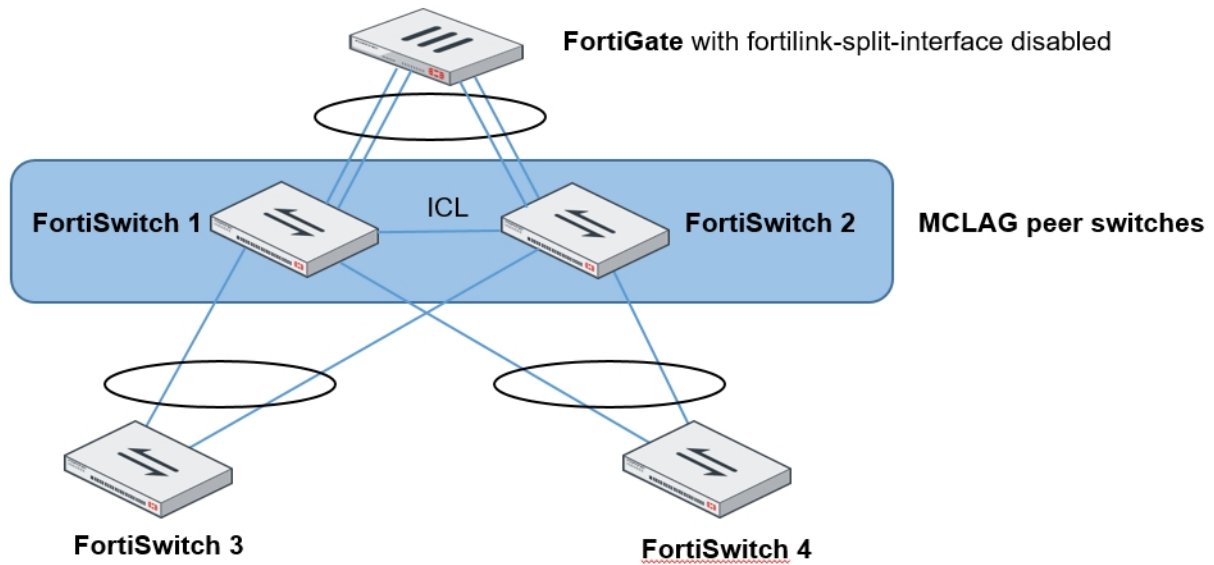
This topology is supported when the FortiGate unit is in HA mode.



Standalone FortiGate unit with dual-homed FortiSwitch access

This network topology provides high port density with two tiers of FortiSwitch units.

Use the `set mclag-icl enable` command to create an ICL on each FortiSwitch unit.

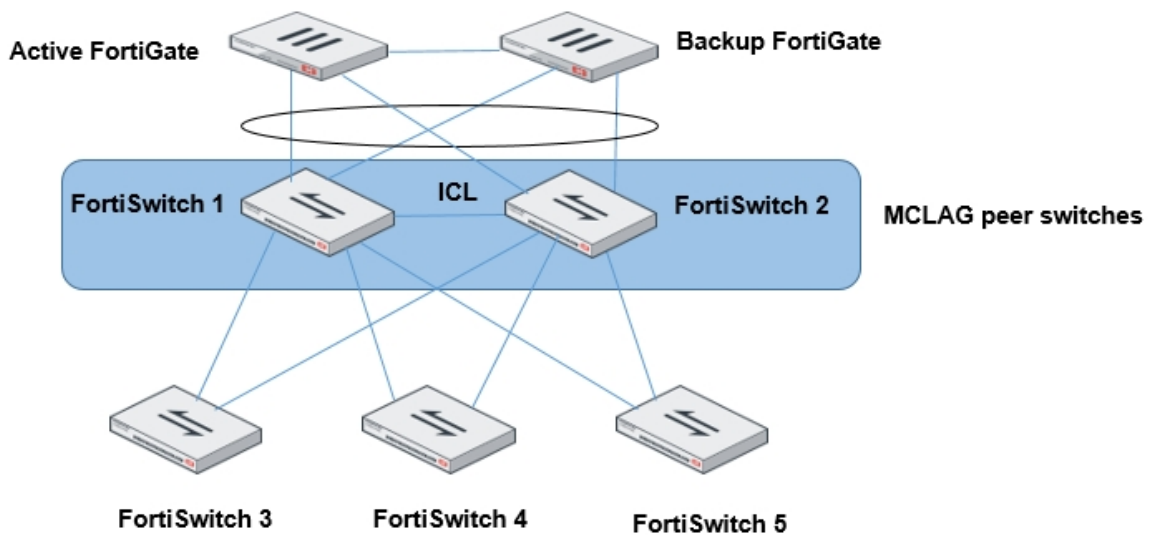


HA-mode FortiGate units with dual-homed FortiSwitch access

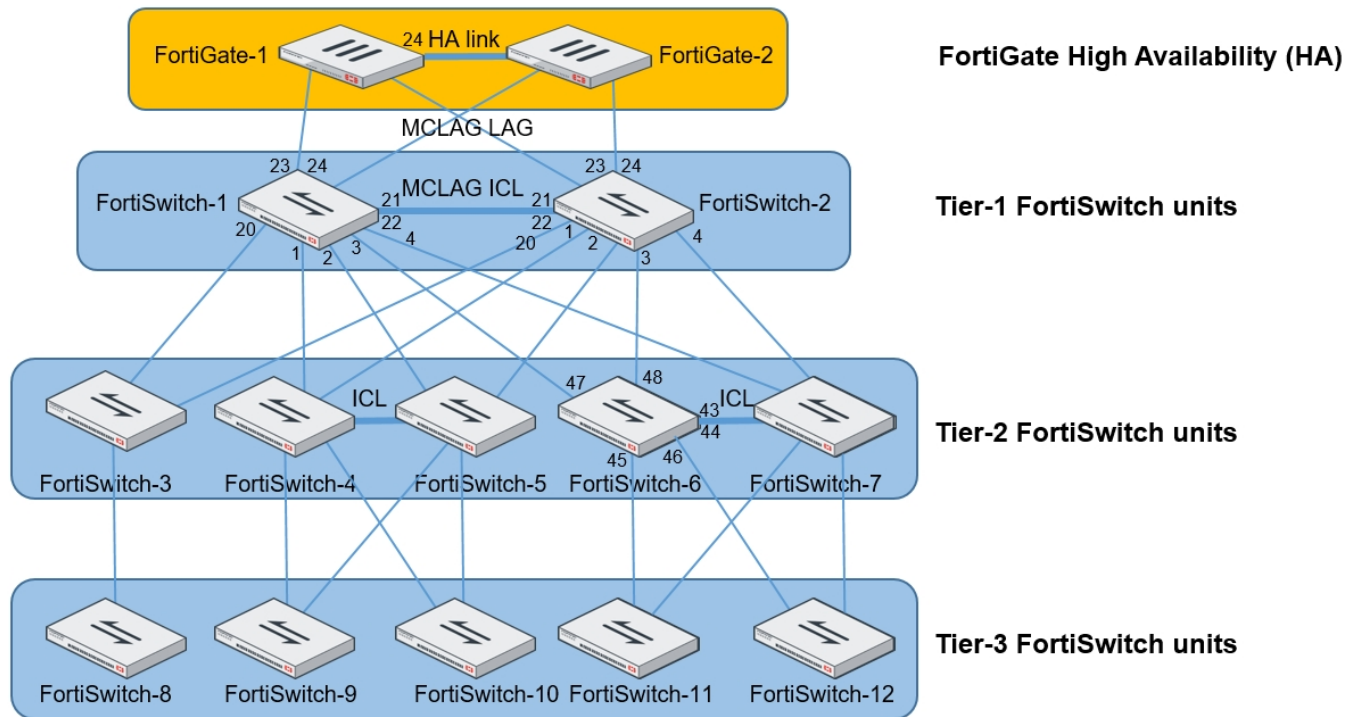
In HA mode, only one FortiGate is active at a time. If the active FortiGate unit fails, the backup FortiGate unit becomes active.

Use the `set mclag-icl enable` command to create an ICL on each FortiSwitch unit.

NOTE: When using HA-mode FortiGate units to manage FortiSwitch units, the HA mode must be active-passive.



Multi-tiered MCLAG with HA-mode FortiGate units



NOTE:

- When using HA-mode FortiGate units to manage FortiSwitch units, the HA mode must be active-passive.
- In this topology, you must use the `auto-isl-port-group` setting as described in the following configuration example. This setting instructs the switches to group ports from MCLAG peers together into one MCLAG when the inter-switch link (ISL) is formed.
- The inter-chassis link (ICL) and `auto-isl-port-group` settings must be done directly on the FortiSwitch unit.
- CLI commands in red are manually configured.
- In a two-tier MCLAG topology, disable STP on the first tier MCLAG peer group with the following commands on each peer switch and do not use access-ring connections on first tier MCLAG peer groups:

```
config switch global
    set mclag-stp-aware disable
end
```

To configure a multi-tiered MCLAG with HA-mode FortiGate units:

1. Configure FortiSwitch-1 for the tier-1 MCLAG:
 - a. Enable the ICL on the ISL formed with the MCLAG peer switch:

```
config switch trunk
    edit "D243Z14000288-0" // trunk name derived from FortiSwitch-2 SN
        set mode lacp-active
        set auto-isl 1
        set mclag-icl enable
```

```
        set members "port21" "port22"
    end
```

- b. Configure the two auto-isl-port-groups based on the topology diagram. The group name must match the name that is configured on the peer switch.

```
config switch auto-isl-port-group
edit "mclag-core1"
    set members "port1" "port2"
next
edit "mclag-core2"
    set members "port3" "port4"
end
```

- c. After you complete the CLI commands in Steps 1a and 1b, the trunks are automatically formed:

```
config switch trunk
edit "D243Z14000288-0"
    set mode lacp-active
    set auto-isl 1
    set mclag-icl enable
    set members "port21" "port22"
next
edit "__FoRtIiLiNk0__"
    set mclag enable
    set members "port24" "port23"
next
edit "8DN4K16000360-0" // trunk name derived from FortiSwitch-3 SN
    set mode lacp-active
    set auto-isl 1
    set mclag enable
    set members "port20"
next
edit "mclag-core1"
    set mode lacp-active
    set auto-isl 1
    set mclag enable
    set members "port1" "port2"
next
edit "mclag-core2"
    set mode lacp-active
    set auto-isl 1
    set mclag enable
    set members "port3" "port4"
next
end
```

2. Configure FortiSwitch-2 for the tier-1 MCLAG:

- a. Enable the ICL on the ISL formed with the MCLAG peer switch:

```
config switch trunk
edit "D243Z14000289-0" // trunk name derived from FortiSwitch-1 SN
    set mode lacp-active
```

```

        set auto-isl 1
        set mclag-icl enable
        set members "port21" "port22"
    end

```

- b. Configure the two auto-isl-port-groups based on the topology diagram. The group name must match the name that is configured on the peer switch.

```

config switch auto-isl-port-group
    edit "mclag-core1"
        set members "port1" "port2"
    next
    edit "mclag-core2"
        set members "port3" "port4"
    end

```

- c. After you complete the CLI commands in Steps 2a and 2b, the trunks are automatically formed:

```

config switch trunk
    edit "D243Z14000288-0"
        set mode lacp-active
        set auto-isl 1
        set mclag-icl enable
        set members "port21" "port22"
    next
    edit "__FoRtI1LiNk0__"
        set mclag enable
        set members "port24" "port23"
    next
    edit "8DN4K16000360-0" // trunk name derived from FortiSwitch-3 SN
        set mode lacp-active
        set auto-isl 1
        set mclag enable
        set members "port20"
    next
    edit "mclag-core1"
        set mode lacp-active
        set auto-isl 1
        set mclag enable
        set members "port1" "port2"
    next
    edit "mclag-core2"
        set mode lacp-active
        set auto-isl 1
        set mclag enable
        set members "port3" "port4"
    next
end

```

3. Tier-2 MLAGs. Enable the ICL between the MLAG peers. For example, configure FortiSwitch-6 as follows.
 - a. Change the tier-2 MLAG peer switches to FortiLink mode and connect them to each other. Enable the ICL on the ISL formed with the MLAG peer switches.

```
config switch trunk
    edit "8DN3X15000026-0" // trunk name derived from FortiSwitch-7 SN
        set mode lacp-active
        set auto-isl 1
        set mclag-icl enable
        set members "port43" "port44"
    end
```

b. The trunks are automatically formed as below:

```
config switch trunk
    edit "8DN3X15000026-0"
        set mode lacp-active
        set auto-isl 1
        set mclag-icl enable
        set members "port43" "port44"
    next
    edit "8EP3X17000051-0" // trunk name derived from FortiSwitch-11 SN
        set mode lacp-active
        set auto-isl 1
        set mclag enable
        set members "port45"
    next
    edit "_FlInK1_MLAG0_"
        set mode lacp-active
        set auto-isl 1
        set mclag enable
        set members "port48" "port47"
    next
    edit "8EP3X17000069-0" // trunk name derived from FortiSwitch-12 SN
        set mode lacp-active
        set auto-isl 1
        set mclag enable
        set members "port46"
    next
end
```

4. Access FortiSwitch units. The access switch trunks are formed automatically as below.

On FortiSwitch-11:

```
config switch trunk
    edit "_FlInK1_MLAG0_"
        set mode lacp-active
        set auto-isl 1
        set mclag enable
        set members "port48" "port47"
    next
end
```

On FortiSwitch-12:

```
config switch trunk
    edit "_FlInK1_MLAG0_"
        set mode lacp-active
        set auto-isl 1
        set mclag enable
        set members "port47" "port48"
    next
end
```


Grouping FortiSwitch units

You can simplify the configuration and management of complex topologies by creating FortiSwitch groups. A group can include one or more FortiSwitch units and you can include different models in a group.

```
config switch-controller switch-group
  edit <name>
    set description <string>
    set members <serial-number> <serial-number> ...
  end
end
```

Grouping FortiSwitch units allows you to restart all of the switches in the group instead of individually. For example, you can use the following command to restart all of the FortiSwitch units in a group named `my-switch-group`:

```
execute switch-controller restart-swtp my-switch-group
```

Upgrading the firmware of FortiSwitch groups is easier, too, because fewer commands are needed. See [Firmware upgrade of stacked or tiered FortiSwitch units on page 1947](#).

Stacking configuration

To set up stacking:

1. Configure the active FortiLink interface on the FortiGate unit.
2. (Optional) Configure the standby FortiLink interface.
3. Connect the FortiSwitch units together, based on your chosen topology.

1. Configure the active FortiLink

Configure the FortiLink interface (as described in the [FortiLink configuration using the FortiGate GUI](#) chapter).

When you configure the FortiLink interface, the stacking capability is enabled automatically.

2. Configure the standby FortiLink

Configure the standby FortiLink interface. Depending on your configuration, the standby FortiLink might connect to the same FortiGate unit as the active FortiLink or to a different FortiGate unit.

If the FortiGate unit receives discovery requests from two FortiSwitch units, the link from one FortiSwitch unit will be selected as active, and the link from other FortiSwitch unit will be selected as standby.

If the active FortiLink fails, the FortiGate unit converts the standby FortiLink to active.

3. Connect the FortiSwitch units

Refer to the topology diagrams to see how to connect the FortiSwitch units.

Inter-switch links (ISLs) form automatically between the stacked switches.


The FortiGate unit will discover and authorize all of the FortiSwitch units that are connected. After this, the FortiGate unit is ready to manage all of the authorized FortiSwitch units.

Disable stacking

To disable stacking, execute the following commands from the FortiGate CLI. In the following example, port4 is the FortiLink interface:

```
config system interface
  edit port4
    set fortilink-stacking disable
  end
end
```

Firmware upgrade of stacked or tiered FortiSwitch units



In this topology, the core FortiSwitch units are model FS-224D-FPOE, and the access FortiSwitch units are model FS-124D-POE. Because the switches are stacked or tiered, the procedure to update the firmware is simpler. In the following procedure, the four FortiSwitch units are upgraded from 3.6.1 to 3.6.2.

To upgrade the firmware of stacked or tiered FortiSwitch units:

- 1. Check that all of the FortiSwitch units are connected and which firmware versions they are running. For example:

```
execute switch-controller get-conn-status

STACK-NAME: FortiSwitch-Stack-port2
SWITCH-ID      VERSION  STATUS      ADDRESS      JOIN-TIME
NAME
S108DV2EJZDAC42F v3.6.0   Authorized/Up 169.254.2.4   Thu Feb  8 17:07:35 2018
-
S108DV4FQON40Q07 v3.6.0   Authorized/Up 169.254.2.5   Thu Feb  8 17:08:37 2018
-
S108DVBWVLH4QGEB v3.6.0   Authorized/Up 169.254.2.6   Thu Feb  8 17:09:13 2018
-
S108DVCY19SA0CD8 v3.6.0   Authorized/Up 169.254.2.2   Thu Feb  8 17:04:41 2018
-
S108DVD98KMQGC44* v3.6.0   Authorized/Up 169.254.2.7   Thu Feb  8 17:10:50 2018
-
S108DVGGBJLQQO48* v3.6.0   Authorized/Up 169.254.2.3   Thu Feb  8 17:06:57 2018
-
```

```

S108DVKM5T2QEA92   v3.6.0   Authorized/Up   169.254.2.8   Thu Feb  8 17:11:00 2018
-
S108DVZX3VTA0045   v3.6.0   Authorized/Up   169.254.2.9   Thu Feb  8 17:11:00 2018
-

Managed-Switches:  8      UP:  8      DOWN:  0

```

2. Upload the firmware image for each FortiSwitch model (FS-224D-FPOE and FS-124D-POE) from either an FTP or TFTP server. If you are using a virtual domain (VDOM), you must enter the `config global` command before entering the `upload-swtp-image` command. For example:

```

FG100E4Q16004478 (global) # execute switch-controller upload-swtp-image tftp FSW_124D_POE-
v3-build0382-FORTINET.out 172.30.12.18

Downloading file FSW_124D_POE-v3-build0382-FORTINET.out from tftp server 172.30.12.18...
#####
Image checking ...
Image MD5 calculating ...
Image Saving S124DP-IMG.swtp ...
Successful!

File Syncing...

FG100E4Q16004478 (global) # execute switch-controller upload-swtp-image tftp FSW_224D_FPOE-
v3-build0382-FORTINET.out 172.30.12.18

Downloading file FSW_224D_FPOE-v3-build0382-FORTINET.out from tftp server 172.30.12.18...
#####
Image checking ...
Image MD5 calculating ...
Image Saving S224DF-IMG.swtp ...
Successful!

File Syncing...

```

3. Check which firmware images are available. For example:

```

FG100E4Q16004478 (root) # execute switch-controller list-swtp-image
SWTP Images on AC:

```

ImageName	ImageSize(B)	ImageInfo	ImageMTime
S124DP-IMG.swtp	19174985	S124DP-v3.6-build382	Mon Oct 2 14:40:54 2017
S224DF-IMG.swtp	23277106	S224DF-v3.6-build382	Mon Oct 2 14:42:55 2017

4. Stage the firmware image for each FortiSwitch model (FS-224D-FPOE and FS-124D-POE). For example:

```

FG100E4Q16004478 (root) # execute switch-controller stage-tiered-swtp-image ALL S124DP-
IMG.swtp
Staged Image Version S124DP-v3.6-build382

FG100E4Q16004478 (root) # execute switch-controller stage-tiered-swtp-image ALL S224DF-
IMG.swtp
Staged Image Version S224DF-v3.6-build382

```

5. Check that the correct firmware image is staged for each FortiSwitch unit. For example:

```

diagnose switch-controller dump network-upgrade status

```

	Running	Status	Next boot

VDOM : root			
S108DVCY19SA0CD8	S108DV-v3.6.0-build4277,171207 (Interim)	(0/0/0)	S108DV-v3.7.0-
build4277,171207	(Interim)		
S108DV2EJZDAC42F	S108DV-v3.6.0-build4277,171207 (Interim)	(0/0/0)	

6. Restart the FortiSwitch units after a 2-minute delay. For example:

```

execute switch-controller restart-swtp-delayed ALL

```

7. When the FortiSwitch units are running again, check that they are running the new firmware version. For example:

```
execute switch-controller get-conn-status
```

```
STACK-NAME: FortiSwitch-Stack-port2
```

SWITCH-ID	VERSION	STATUS	ADDRESS	JOIN-TIME
NAME				
S108DV2EJZDAC42F	v3.6.0	Authorized/Up	169.254.2.4	Thu Feb 8 17:07:35 2018
-				
S108DV4FQON40Q07	v3.6.0	Authorized/Up	169.254.2.5	Thu Feb 8 17:08:37 2018
-				
S108DVBWVLH4QGEB	v3.6.0	Authorized/Up	169.254.2.6	Thu Feb 8 17:09:13 2018
-				
S108DVCY19SA0CD8	v3.6.0	Authorized/Up	169.254.2.2	Thu Feb 8 17:04:41 2018
-				
S108DVD98KMQGC44*	v3.6.0	Authorized/Up	169.254.2.7	Thu Feb 8 17:10:50 2018
-				
S108DVGGBJLQQO48*	v3.6.0	Authorized/Up	169.254.2.3	Thu Feb 8 17:06:57 2018
-				
S108DVKM5T2QEA92	v3.6.0	Authorized/Up	169.254.2.8	Thu Feb 8 17:11:00 2018
-				
S108DVZX3VTA0045	v3.6.0	Authorized/Up	169.254.2.9	Thu Feb 8 17:11:00 2018
-				

```
Managed-Switches: 8      UP: 8      DOWN: 0
```

Transitioning from a FortiLink split interface to a FortiLink MCLAG

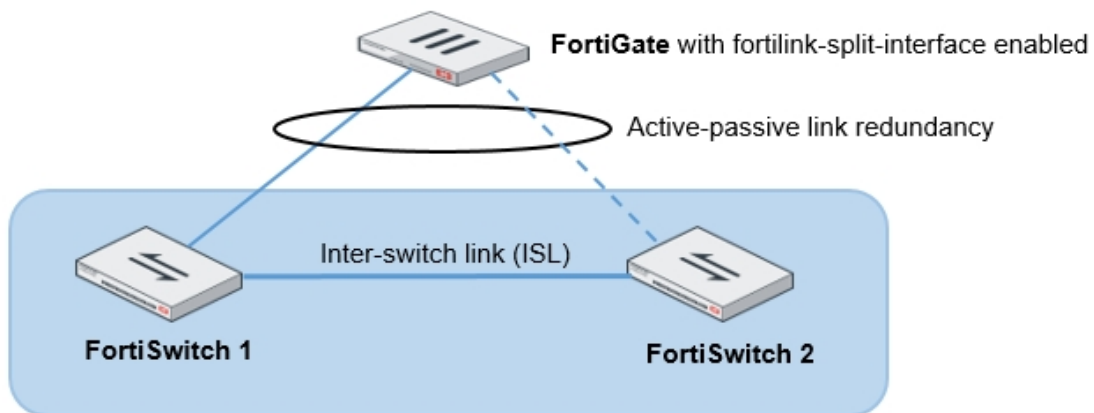
In this topology, the FortiLink split interface connects a FortiLink aggregate interface from one FortiGate unit to two FortiSwitch units.

NOTE:

- This procedure also applies to a FortiGate unit in HA mode.
- More links can be added between the FortiGate unit and FortiSwitch unit.
- After the MCLAG is set up, only connect the tier-2 FortiSwitch units.
- When you are using the aggregate interface on the FortiGate unit for the FortiLink interface, the `lacp-mode` of the FortiLink aggregate interface must be set to `static`.

1. Enable the split interface on the FortiLink aggregate interface. By default, the split interface is enabled. For example:

```
config system interface
  edit flinksplit1
    set ip 169.254.3.1 255.255.255.0
    set allowaccess ping capwap https
    set vlanforward enable
    set type aggregate
    set member port4 port5
    set lacp-mode static
    set fortilink enable
    set fortilink-split-interface enable
  next
end
```



2. Log into FortiSwitch 2 using the **Connect to CLI** button in the FortiGate GUI, use the `get switch lldp auto-isl-status` command to find out the name of the trunk connecting the peer switches, and change the ISL to an ICL. For example:

```
get switch lldp auto-isl-status

config switch trunk
  edit <trunk_name>
    set mclag-icl enable
```

```

    next
end

```

3. Log into FortiSwitch 1 using the *Connect to CLI* button in the FortiGate GUI, use the `get switch lldp auto-isl-status` command to find out the name of the trunk connecting the peer switches, and change the ISL to an ICL. For example:

```

get switch lldp auto-isl-status

config switch trunk
    edit <trunk_name>
        set mclag-icl enable
    next
end

```

4. Log into the FortiGate unit and disable the split interface. For example:

```

config system interface
    edit flinksplit1
        set fortilink-split-interface disable
    next
end

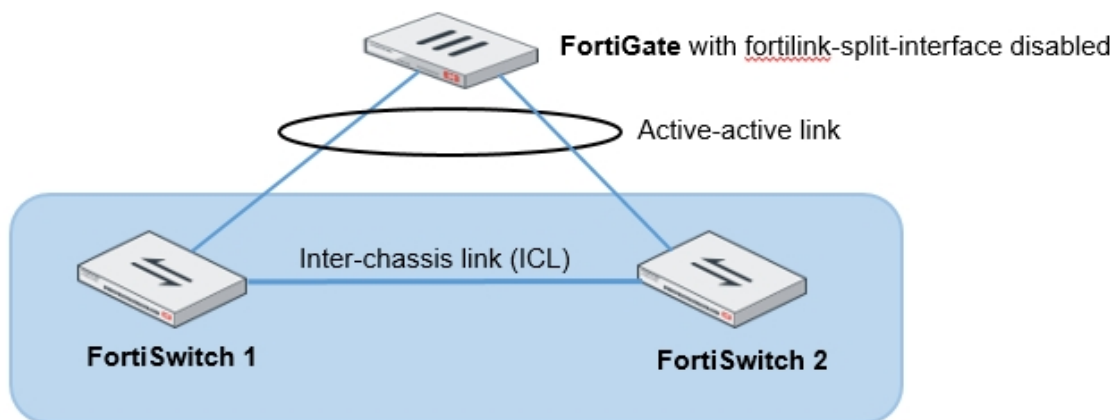
```

5. Enable the LACP active mode.
6. Check that the LAG is working correctly. For example:

```

diagnose netlink aggregate name <aggregate_name>

```



Optional setup tasks

This section describes the following tasks:

- [Configuring the FortiSwitch management port on page 1952](#)
- [Converting to FortiSwitch standalone mode on page 1953](#)
- [Changing the admin password on the FortiGate for all managed FortiSwitch units on page 1953](#)
- [Enabling network-assisted device detection on page 1954](#)
- [Limiting the number of parallel process for FortiSwitch configuration on page 1954](#)

Configuring the FortiSwitch management port

If the FortiSwitch model has a dedicated management port, you can configure remote management to the FortiSwitch. In FortiLink mode, the FortiGate is the default gateway, so you need to configure an explicit route for the FortiSwitch management port.

Using the Web administration GUI

1. Go to *Network > Static Routes > Create New > Route*.
2. Set *Destination* to *Subnet* and enter a subnetwork and mask.
3. Set *Device* to the management interface.
4. Add a *Gateway* IP address.

Using the FortiSwitch CLI

Enter the following commands:

```
config router static
  edit 1
    set device mgmt
    set gateway <router IP address>
    set dst <router subnet> <subnet mask>
  end
end
```

In the following example, the FortiSwitch management port is connected to a router with IP address 192.168.0.10:

```
config router static
  edit 1
    set device mgmt
    set gateway 192.168.0.10
    set dst 192.168.0.0 255.255.0.0
  end
end
```

Converting to FortiSwitch standalone mode

Use one of the following commands to convert a FortiSwitch from FortiLink mode to standalone mode so that it will no longer be managed by a FortiGate:

- `execute switch-controller factory-reset <switch-id>`
This command returns the FortiSwitch to the factory defaults and then reboots the FortiSwitch. If the FortiSwitch is configured for FortiLink auto-discovery, FortiGate can detect and automatically authorize the FortiSwitch. For example:
`execute switch-controller factory-reset S1234567890`
- `execute switch-controller set-standalone <switch-id>`
This command returns the FortiSwitch to the factory defaults, reboots the FortiSwitch, and prevents the FortiGate from automatically detecting and authorizing the FortiSwitch. For example:
`execute switch-controller set-standalone S1234567890`

You can disable FortiLink auto-discovery on multiple FortiSwitch units using the following commands:

```
config switch-controller global
    set disable-discovery <switch-id>
end
```

For example:

```
config switch-controller global
    set disable-discovery S1234567890
end
```

You can also add or remove entries from the list of FortiSwitch units that have FortiLink auto-discovery disabled using the following commands:

```
config switch-controller global
    append disable-discovery <switch-id>
    unselect disable-discovery <switch-id>
end
```

For example:

```
config switch-controller global
    append disable-discovery S012345678
    unselect disable-discovery S1234567890
end
```

Changing the admin password on the FortiGate for all managed FortiSwitch units

By default, each FortiSwitch has an admin account without a password. To replace the admin passwords for all FortiSwitch units managed by a FortiGate, use the following commands from the FortiGate CLI:

```
config switch-controller switch-profile
    edit default
        set login-passwd-override {enable | disable}
        set login-passwd <password>
    next
end
```


If you had already applied a profile with the override enabled and the password set and then decide to remove the admin password, you need to apply a profile with the override enabled and no password set; otherwise, your previously set password will remain in the FortiSwitch. For example:

```
config switch-controller switch-profile
  edit default
    set login-passwd-override enable
    unset login-passwd
  next
end
```

Enabling network-assisted device detection

Network-assisted device detection allows the FortiGate unit to use the information about connected devices detected by the managed FortiSwitch unit.

To enable network-assisted device detection on a VDOM:

```
config switch-controller network-monitor-settings
  set network-monitoring enable
end
```

You can display a list of detected devices from the *Device Inventory* menu in the GUI. To list the detected devices in the CLI, enter the following command:

```
diagnose user device list
```

Limiting the number of parallel process for FortiSwitch configuration

Use the following CLI commands to reduce the number of parallel process that the switch controller uses for configuring FortiSwitch units:

```
config global
  config switch-controller system
    set parallel-process-override enable
    set parallel-process <1-300>
  end
end
```

FortiSwitch features configuration

This section describes how to configure global FortiSwitch settings using FortiGate CLI commands. These settings will apply to all of the managed FortiSwitch units. You can also override some of the settings on individual FortiSwitch units.

This chapter covers the following topics:

- ["Configure VLANs" on page 1955](#)
- ["Configure IGMP settings " on page 1958](#)
- ["Configure LLDP-MED" on page 1958](#)
- ["Configure the MAC sync interval " on page 1960](#)
- ["Configure STP settings" on page 1960](#)
- ["Quarantines" on page 1961](#)

Configure VLANs

Use Virtual Local Area Networks (VLANs) to logically separate a LAN into smaller broadcast domains. VLANs allow you to define different policies for different types of users and to set finer control on the LAN traffic. (Traffic is only sent automatically within the VLAN. You must configure routing for traffic between VLANs.)

From the FortiGate unit, you can centrally configure and manage VLANs for the managed FortiSwitch units.

In FortiSwitchOS 3.3.0 and later releases, the FortiSwitch supports untagged and tagged frames in FortiLink mode. The switch supports up to 1,023 user-defined VLANs. You can assign a VLAN number (ranging from 1-4095) to each of the VLANs.

You can configure the default VLAN for each FortiSwitch port as well as a set of allowed VLANs for each FortiSwitch port.

FortiSwitch VLANs display

The *WiFi & Switch Controller > FortiSwitch VLANs* page displays VLAN information for the managed switches.

<div> + Create New Edit Delete Q Search </div>				
Name	VLAN ID	IP/Netmask	Access	Ref.
vlan44	44	192.168.2.1 255.255.255.0	SNMP	0
vlan45	45	10.10.10.1 255.255.255.0		1
vsw.port3	1	172.20.20.10 255.255.255.0	HTTPS HTTP	10

Each entry in the VLAN list displays the following information:

- *Name*—name of the VLAN
- *VLAN ID*—the VLAN number
- *IP/Netmask*—address and mask of the subnetwork that corresponds to this VLAN
- *Access*—administrative access settings for the VLAN
- *Ref*—number of configuration objects referencing this VLAN

Enabling and disabling switch-controller access VLANs through the FortiGate unit

Access VLANs are VLANs that aggregate client traffic solely to the FortiGate unit. This prevents direct client-to-client traffic visibility at the layer-2 VLAN layer. Clients can only communicate with the FortiGate unit. After the client traffic reaches the FortiGate, the FortiGate unit can then determine whether to allow various levels of access to the client by shifting the client's network VLAN as appropriate.

NOTE: IPv6 is not supported between clients within a switch-controller access VLAN.

Use `enable` to allow traffic only to and from the FortiGate and to block FortiSwitch port-to-port traffic on the specified VLAN. Use `disable` to allow normal traffic on the specified VLAN.

```
config system interface
    edit <VLAN name>
        set switch-controller-access-vlan {enable | disable}
    next
end
```

NOTE: You must configure the proxy ARP with the `config system proxy-arp` CLI command to be able to use the access VLANs. For example:

```
config system proxy-arp
  edit 1
    set interface "V100"
    set ip 1.1.1.1
    set end-ip 1.1.1.200
  next
end
```

Creating VLANs

Setting up a VLAN requires you to create the VLAN and assign FortiSwitch ports to the VLAN. You can do this with either the Web GUI or CLI.

Using the Web administration GUI

To create the VLAN:



1. Go to *WiFi & Switch Controller > FortiSwitch VLANs*, select *Create New*, and change the following settings:

Interface Name	VLAN name
VLAN ID	Enter a number (1-4094)
Color	Choose a unique color for each VLAN, for ease of visual display.
IP/Network Mask	IP address and network mask for this VLAN.

2. Enable *DHCP Server* and set the IP range.
3. Set the *Admission Control* options as required.
4. Select *OK*.

To assign FortiSwitch ports to the VLAN:

1. Go to *WiFi & Switch Controller > FortiSwitch Ports*.
2. Click the desired port row.
3. Click the *Native VLAN* column in one of the selected entries to change the native VLAN.
4. Select a VLAN from the displayed list. The new value is assigned to the selected ports.
5. Click the + icon in the *Allowed VLANs* column to change the allowed VLANs.
6. Select one or more of the VLANs (or the value *all*) from the displayed list. The new value is assigned to the selected port.

Port	Description	Native VLAN	Allowed VLANs	Device Information	PoE	Bytes (Sent/Received)
My-Switch - FS108D3W16001177 (10)						
<div>   <div>PoE Status:</div> <div>Total Budget: 75.0W</div> </div>						
port1		vsw.port3			Powered	0 B
port2		vsw.port3			Powered	0 B
port3		vlan45			Powered	0 B
port4		vlan45			Powered	0 B
port5		vlan45			Powered	0 B
port6		vsw.port3	vlan44		Powered	0 B
port7		vsw.port3	vlan44		Powered	0 B
port8		vsw.port3	vlan44		Powered	0 B
port9		vsw.port3	vlan44			0 B
port10		FGVM010000088418				33.27 MB

Using the FortiSwitch CLI

1. Create the marketing VLAN.

```
config system interface
  edit <vlan name>
    set vlanid <1-4094>
    set color <1-32>
    set interface <FortiLink-enabled interface>
  end
```

2. Set the VLAN's IP address.

```
config system interface
  edit <vlan name>
    set ip <IP address> <Network mask>
  end
```

3. Enable a DHCP Server.

```
config system dhcp server
  edit 1
    set default-gateway <IP address>
    set dns-service default
    set interface <vlan name>
    config ip-range
      set start-ip <IP address>
      set end-ip <IP address>
    end
    set netmask <Network mask>
  end
```

4. Assign ports to the VLAN.

```
config switch-controller managed-switch
  edit <Switch ID>
```

```
config ports
  edit <port name>
    set vlan <vlan name>
    set allowed-vlans <vlan name>
    or
    set allowed-vlans-all enable
  next
end
end
```

Assign untagged VLANs to a managed FortiSwitch port:

```
config switch-controller managed-switch
  edit <managed-switch>
    config ports
      edit <port>
        set untagged-vlans <VLAN-name>
      next
    end
  next
end
```

Configure IGMP settings

Use the following command to configure the global IGMP settings.

Aging time is the maximum number of seconds that the system will retain a multicast snooping entry. Enter an integer value from 15 to 3600. The default value is 300.

Flood-unknown-multicast controls whether the system will flood unknown multicast messages within the VLAN.

```
config switch-controller igmp-snooping
  set aging-time <15-3600>
  set flood-unknown-multicast {enable | disable}
end
```

Configure LLDP-MED

To configure LLDP profiles:

```
config switch-controller lldp-profile
  edit <profile number>
    set 802.1-tlvs port-vlan-id
    set 802.3-tlvs max-frame-size
    set auto-isl {enable | disable}
    set auto-isl-hello-timer <1-30>
    set auto-isl-port-group <0-9>
    set auto-isl-receive-timeout <3-90>
    set med-tlvs (inventory-management | network-policy)
  end
```

To configure LLDP settings:

```
config switch-controller lldp-settings
  set status < enable | disable >
```

```

set tx-hold <int>
set tx-interval <int>
set fast-start-interval <int>
set management-interface {internal | management}
end

```

Variable	Description
status	Enable or disable
tx-hold	Number of tx-intervals before the local LLDP data expires. Therefore, the packet TTL (in seconds) is tx-hold times tx-interval . The range for tx-hold is 1 to 16, and the default value is 4.
tx-interval	How often the FortiSwitch transmits the LLDP PDU. The range is 5 to 4095 seconds, and the default is 30 seconds.
fast-start-interval	How often the FortiSwitch transmits the first 4 LLDP packets when a link comes up. The range is 2 to 5 seconds, and the default is 2 seconds. Set this variable to zero to disable fast start.
management-interface	Primary management interface to be advertised in LLDP and CDP PDUs.

Create LLDP asset tags for each managed FortiSwitch

You can use the following commands to add an LLDP asset tag for a managed FortiSwitch:

```

config switch-controller managed-switch
edit <fsw>
set switch-device-tag <string>
end

```

Add media endpoint discovery (MED) to an LLDP configuration

You can use the following commands to add media endpoint discovery (MED) features to an LLDP profile:

```

config switch-controller lldp-profile
edit <lldp-profile>
config med-network-policy
edit guest-voice
set status {disable | enable}
next
edit guest-voice-signaling
set status {disable | enable}
next
edit guest-voice-signaling
set status {disable | enable}
next
edit softphone-voice
set status {disable | enable}
next
edit streaming-video
set status {disable | enable}
next
edit video-conferencing
set status {disable | enable}
end

```

```
    next
    edit video-signaling
        set status {disable | enable}
    next
    edit voice
        set status {disable | enable}
    next
    edit voice-signaling
        set status {disable | enable}
    end
config custom-tlvs
    edit <name>
        set oui <identifier>
        set subtype <subtype>
        set information-string <string>
    end
end
```

Display LLDP information

You can use the following commands to display LLDP information:

```
diagnose switch-controller dump lldp stats <switch> <port>
diagnose switch-controller dump lldp neighbors-summary <switch>
diagnose switch-controller dump lldp neighbors-detail <switch>
```

Configure the MAC sync interval

Use the following commands to configure the global MAC sync interval.

The MAC sync interval is the time interval between MAC synchronizations. The range is 30 to 600 seconds, and the default value is 60.

```
config switch-controller mac-sync-settings
    set mac-sync-interval <30-600>
end
```

Configure STP settings

NOTE: STP is not supported between a FortiGate unit and a FortiSwitch unit in FortiLink mode.

Use the following CLI commands for global STP configuration. This configuration applies to all managed FortiSwitch units:

```
config switch-controller stp-settings
    set name <name>
    set revision <stp revision>
    set hello-time <hello time>
    set forward-time <forwarding delay>
    set max-age <maximum aging time>
    set max-hops <maximum number of hops>
end
```

You can override the global STP settings for a FortiSwitch unit using the following commands:

```
config switch-controller managed-switch
```

```
edit <switch-id>
  config stp-settings
    set local-override enable
```

Quarantines

Administrators can use MAC addresses to quarantine hosts and users connected to a FortiSwitch unit. Quarantined MAC addresses are isolated from the rest of the network and LAN by using a separate VLAN.

Quarantining MAC addresses

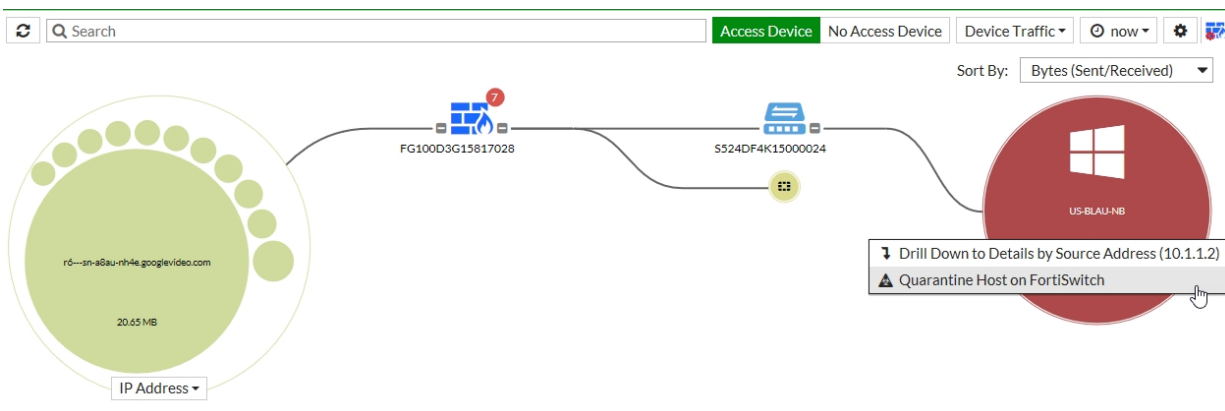
You can use the FortiGate GUI or CLI to quarantine a MAC address.

NOTE: If you have multiple FortiLink interfaces, only the first quarantine VLAN is created successfully (with an IP address of 10.254.254.254). Additional quarantine VLANs will have an empty IP address.

Using the FortiGate GUI

In the FortiGate GUI, the quarantine feature is automatically enabled when you quarantine a host.

1. Select the host to quarantine.
 - Go to *Security Fabric > Physical Topology*, right-click on a host, and select *Quarantine Host on FortiSwitch*.
 - Go to *Security Fabric > Logical Topology*, right-click on a host, and select *Quarantine Host on FortiSwitch*.
 - Go to *FortiView > Sources*, right-click on an entry in the Source column, and select *Quarantine Host on FortiSwitch*.
2. Select *Accept* to confirm that you want to quarantine the host.



Using the FortiGate CLI

NOTE: Previously, this feature used the `config switch-controller quarantine` CLI command.

By default, the quarantine feature is enabled. When you upgrade a FortiGate unit from an older to a newer firmware version, the FortiGate unit uses the quarantine feature status from the older configuration. If the quarantine feature was disabled in the older configuration, it will be disabled after the upgrade.

You can add MAC addresses to be quarantined even when the quarantine feature is disabled. The MAC addresses are only quarantined when the quarantine feature is enabled.

The table size limit for the quarantine entry is 512. There is no limit for how many MAC addresses can be quarantined per quarantine entry.

```
config user quarantine
  set quarantine enable
  config targets
    edit <quarantine_entry_name>
      set description <string>
      config macs
        edit <MAC_address_1>
        next
        edit <MAC_address_2>
        next
        edit <MAC_address_3>
        next
      end
    end
  end
```

Option	Description
quarantine_entry_name	A name for this quarantine entry.
string	Optional. A description of the MAC addresses being quarantined.
MAC_address_1, MAC_address_2, MAC_address_3	A layer-2 MAC address in the following format: 12:34:56:aa:bb:cc

For example:

```
config user quarantine
  set quarantine enable
  config targets
    edit quarantine1
      config macs
        set description "infected by virus"
        edit 00:00:00:aa:bb:cc
        next
        edit 00:11:22:33:44:55
        next
        edit 00:01:02:03:04:05
        next
      end
    end
  end
```

Viewing quarantine entries

Quarantine entries are created on the FortiGate unit that is managing the FortiSwitch unit.

Using the FortiGate GUI

1. Go to *Monitor > Quarantine Monitor*.

2. Click *Quarantined on FortiSwitch*.

The Quarantined on FortiSwitch button is only available if a device is detected behind the FortiSwitch unit, which requires Device Detection to be enabled.

Refresh	Delete	Remove All	Search	All	Quarantined on FortiSwitch	Banned IP
Type	Details	Source	Expires	Description		
MAC address	18:db:f2:32:52:e7 (US-BLAU-NB)	Administrative	Never	Hostname: US-BLAU-NB, Use...		

Using the FortiGate CLI

Use the following command to view the quarantine list of MAC addresses:

```
show user quarantine
```

For example:

```
show user quarantine

config user quarantine
  set quarantine enable
  config targets
    edit quarantine1
      config macs
        set description "infected by virus"
        edit 00:00:00:aa:bb:cc
        next
        edit 00:11:22:33:44:55
        next
        edit 00:01:02:03:04:05
        next
      end
    end
  end
end
```

When the quarantine feature is enabled on the FortiGate unit, it creates a quarantine VLAN (qtn.<FortiLink_port_name>) and a quarantine DHCP server (with the quarantine VLAN as default gateway) on the virtual domain. The quarantine VLAN is applied to the allowed and untagged VLANs on all connected FortiSwitch ports.

Use the following command to view the quarantine VLAN:

```
show system interface qtn.<FortiLink_port_name>
```

For example:

```
show system interface qtn.port7

config system interface
  edit "qtn.port7"
    set vdom "vdom1"
    set ip 10.254.254.254 255.255.255.0
    set description "Quarantine VLAN"
    set security-mode captive-portal
    set replacemsg-override-group "auth-intf-qtn.port7"
    set device-identification enable
    set device-identification-active-scan enable
    set snmp-index 34
    set switch-controller-access-vlan enable
    set color 6
```

```

        set interface "port7"
        set vlanid 4093
    next
end

```

Use the following commands to view the quarantine DHCP server:

```

show system dhcp server
config system dhcp server
    edit 2
        set dns-service default
        set default-gateway 10.254.254.254
        set netmask 255.255.255.0
        set interface "qtn.port7"
        config ip-range
            edit 1
                set start-ip 10.254.254.192
                set end-ip 10.254.254.253
            next
        end
        set timezone-option default
    next
end

```

Use the following command to view how the quarantine VLAN is applied to the allowed and untagged VLANs on all connected FortiSwitch ports:

```

show switch-controller managed-switch

```

For example:

```

show switch-controller managed-switch







config switch-controller managed-switch
    edit "FS1D483Z15000036"
        set fsw-wan1-peer "port7"
        set fsw-wan1-admin enable
        set version 1
        set dynamic-capability 503
        config ports
            edit "port1"
                set vlan "vsw.port7"
                set allowed-vlans "qtn.port7"
                set untagged-vlans "qtn.port7"
            next
            edit "port2"
                set vlan "vsw.port7"
                set allowed-vlans "qtn.port7"
                set untagged-vlans "qtn.port7"
            next
            edit "port3"
                set vlan "vsw.port7"
                set allowed-vlans "qtn.port7"
                set untagged-vlans "qtn.port7"
            next
            ...
        end
    end
end

```

Releasing MAC addresses from quarantine

Using the FortiGate GUI

1. Go to *Monitor > Quarantine Monitor*.
2. Click *Quarantined on FortiSwitch*.
3. Right-click on one of the entries and select *Delete* or *Remove All*.
4. Click *OK* to confirm your choice.

 Refresh		 Delete	 Remove All	 Search	All	Quarantined on FortiSwitch	Banned IP
▼ Type		▼ Details		▼ Source	▼ Expires	▼ Description	
MAC address		18:00:00:00:00:00 (US-BLAU-NB)		Administrative	Never	Hostname: US-BLAU-NB, Use...	
		<div><div> Delete</div><div> Remove All</div></div>					

Using the FortiGate CLI

To release MAC addresses from quarantine, you can delete a single MAC address or delete a quarantine entry, which will delete all of the MAC addresses listed in the entry. You can also disable the quarantine feature, which releases all quarantined MAC addresses from quarantine.

To delete a single quarantined MAC address:

```
config user quarantine
  config targets
    edit <quarantine_entry_name>
      config macs
        delete <MAC_address_1>
      end
    end
  end
end
```

To delete all MAC addresses in a quarantine entry:

```
config user quarantine
  config targets
    delete <quarantine_entry_name>
  end
end
```

To disable the quarantine feature:

```
config user quarantine
  set quarantine disable
end
```

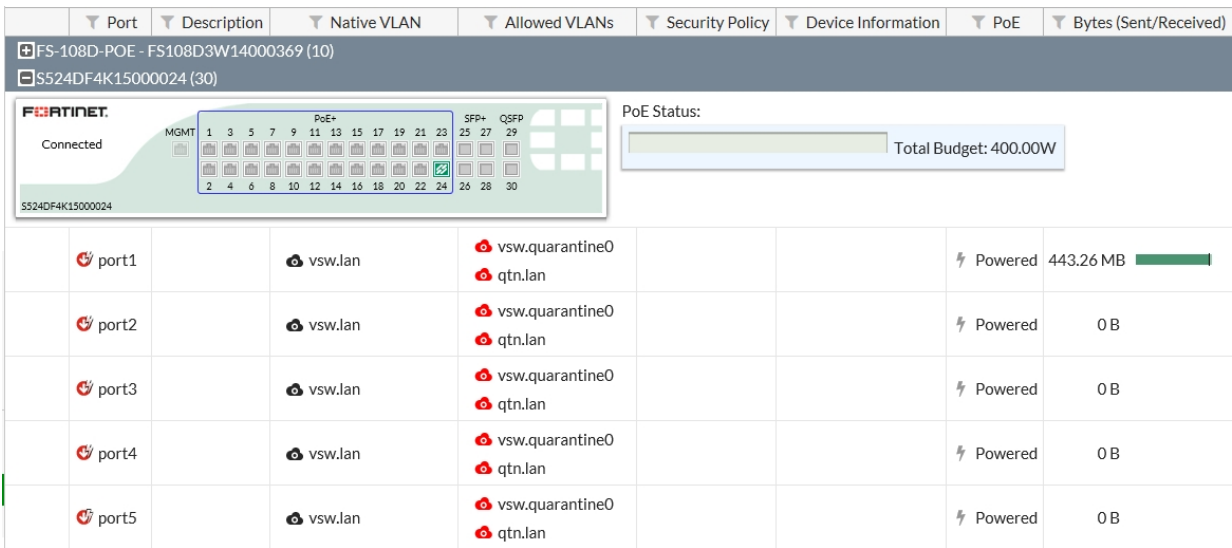
FortiSwitch port features

You can configure the FortiSwitch port feature settings from the FortiGate using the FortiSwitch CLI or web administration GUI.

FortiSwitch ports display

The *WiFi & Switch Controller > FortiSwitch Ports* page displays port information about each of the managed switches.

The following figure shows the display for a FortiSwitch 524D-FPOE:




The switch faceplate displays:

- active ports (green)
- PoE-enabled ports (blue rectangle)
- FortiLink port (link icon)

PoE Status displays the total power budget and the actual power currently allocated.

The allocated power displays a blue bar for the used power (currently being consumed) and a green bar for the reserved power (power available for additional devices on the POE ports). See the following figures:



Each entry in the port list displays the following information:

- Port status (red for down, green for up)
- Port name

- Native VLAN
- Allowed VLANs
- Device information
- PoE status
- Bytes sent and received by the port

Configuring ports using the GUI

You can use the *WiFi & Switch Controller > FortiSwitch Ports* page to do the following with FortiSwitch switch ports:

- Set the native VLAN and add more VLANs
- Edit the description of the port
- Enable or disable the port
- Enable or disable PoE for the port
- Enable or disable DHCP blocking (if supported by the port)
- Enable or disable IGMP snooping (if supported by the port)
- Enable or disable whether a port is an edge port
- Enable or disable STP (if supported by the port)
- Enable or disable loop guard (if supported by the port)
- Enable or disable STP BPDU guard (if supported by the port)
- Enable or disable STP root guard (if supported by the port)

Resetting PoE-enabled ports

If you need to reset PoE-enabled ports, go to *WiFi & Switch Control > FortiSwitch Ports*, right-click on one or more PoE-enabled ports and select *Reset PoE* from the context menu.

You can also go to *WiFi & Switch Control > Managed FortiSwitch* and click on a port icon for the FortiSwitch of interest. In the FortiSwitch Ports page, right-click on one or more PoE-enabled ports and select *Reset PoE* from the context menu.

Configuring ports using the FortiGate CLI

You can configure the following FortiSwitch port settings using the FortiGate CLI:

- [Configuring port speed and status on page 1968](#)
- [Configure a VLAN on the port \(see \[Configure VLANs\]\(#\)\)](#)
- [Sharing FortiSwitch ports between VDOMs on page 1968](#)
- [Limiting the number of learned MAC addresses on a FortiSwitch interface on page 1970](#)
- [Configuring the DHCP trust setting on page 1972](#)
- [Configuring PoE on page 1973](#)
- [Configuring edge ports on page 1974](#)
- [Configuring STP on page 1974](#)
- [Configuring STP root guard on page 1975](#)
- [Configuring STP BPDU guard on page 1976](#)
- [Configuring loop guard on page 1978](#)

- [Configuring LLDP settings on page 1978](#)
- [Configuring IGMP settings on page 1979](#)
- [Configuring sFlow on page 1979](#)
- [Configuring Dynamic ARP inspection \(DAI\) on page 1980](#)
- [Configuring FortiSwitch port mirroring on page 1981](#)

Configuring port speed and status

Use the following commands to set port speed and other base port settings:

```
config switch-controller managed-switch
edit <switch>
config ports
edit <port>
set description <text>
set speed <speed>
set status {down | up}
end
end
```

For example:

```
config switch-controller managed-switch
edit S524DF4K15000024
config ports
edit port1
set description "First port"
set speed auto
set status up
end
end
```

Sharing FortiSwitch ports between VDOMs

Virtual domains (VDOMs) are a method of dividing a FortiGate unit into two or more virtual units that function as multiple independent units. VDOMs provide separate security domains that allow separate zones, user authentication, security policies, routing, and VPN configurations.

FortiSwitch ports can now be shared between VDOMs.

NOTE: You cannot use the quarantine feature while sharing FortiSwitch ports between VDOMs.

To share FortiSwitch ports between VDOMs:

1. Create one or more VDOMs.
2. Assign VLANs to each VDOM as required.
3. From these VLANs, select one VLAN to be the default VLAN for the ports in the virtual switch:

```
config switch-controller global
set default-virtual-switch-vlan <VLAN>
```

NOTE: You must execute these commands from the VDOM that the default VLAN belongs to.

When you add a new port to the VDOM, the new port will be automatically assigned to the default VLAN. You can reassign the ports to other VLANs later.

4. Create a virtual port pool (VPP) to contain the ports to be shared:

```
config switch-controller virtual-port-pool
  edit <VPP_name>
    description <string>
  next
end
```

NOTE: You must execute these commands from the VDOM that the default VLAN belongs to.

For example:

```
config switch-controller virtual-port-pool
  edit "pool3"
    description "pool for port3"
  next
end
```

5. Share a FortiSwitch port from the VDOM that the FortiSwitch belongs to with another VDOM or export the FortiSwitch port to a VPP where it can be used by any VDOM:

```
config switch-controller managed-switch
  edit <switch.id>
    config ports
      edit <port_name>
        set {export-to-pool <VPP_name> | export-to <VDOM_name>}
        set export-tags <string1,string2,string3,...>
      next
    end
  next
end
```

NOTE: You must execute these commands from the VDOM that the default VLAN belongs to.

For example, if you want to export a port to the VPP named `pool3`:

```
config switch-controller managed-switch
  edit "S524DF4K15000024"
    config ports
      edit port3
        set export-to-pool "pool3"
        set export-tags "Pool 3"
      next
    end
  next
end
```

For example, if you want to export a port to the VDOM named `vdom3`:

```
config switch-controller managed-switch
  edit "S524DF4K15000024"
    config ports
      edit port3
        set export-to "vdom3"
        set export-tags "VDOM 3"
```



```

        next
    end
next
end

```

6. Request a port in a VPP:

```
execute switch-controller virtual-port-pool request <FortiSwitch_device_ID> <port_name>
```

NOTE: You must execute this command from the VDOM that is requesting the port.

For example:

```
execute switch-controller virtual-port-pool request S524DF4K15000024h port3
```

7. Return a port to a VPP:

```
execute switch-controller virtual-port-pool return <FortiSwitch_device_ID> <port_name>
```

NOTE: You must execute this command from the VDOM that owns the port.

For example:

```
execute switch-controller virtual-port-pool return S524DF4K15000024h port3
```

You can create your own export tags using the following CLI commands:

```

config switch-controller switch-interface-tag
    edit <tag_name>
end

```

Use the following CLI command to list the contents of a specific VPP:

```
execute switch-controller virtual-port-pool show-by-pool <VPP_name>
```

Use the following CLI command to list all VPPs and their contents:

```
execute switch-controller virtual-port-pool show
```

NOTE: Shared ports do not support the following features:

- LLDP
- 802.1x
- STP
- BPDU guard
- Root guard
- DHCP snooping
- IGMP snooping
- QoS
- Port security
- MCLAG

Limiting the number of learned MAC addresses on a FortiSwitch interface

You can limit the number of MAC addresses learned on a FortiSwitch interface (port or VLAN). The limit ranges from 1 to 128. If the limit is set to the default value zero, there is no learning limit.

NOTE: Static MAC addresses are not counted in the limit. The limit refers only to learned MAC addresses.

Use the following CLI commands to limit MAC address learning on a VLAN:

```
config switch vlan
  edit <integer>
    set switch-controller-learning-limit <limit>
  end
end
```

For example:

```
config switch vlan
  edit 100
    set switch-controller-learning-limit 20
  end
end
```

Use the following CLI commands to limit MAC address learning on a port:

```
config switch-controller managed-switch
  edit <FortiSwitch_Serial_Number>
    config ports
      edit <port>
        set learning-limit <limit>
      next
    end
  end
end
```

For example:

```
config switch-controller managed-switch
  edit S524DF4K15000024
    config ports
      edit port3
        set learning-limit 50
      next
    end
  end
end
```

You can change how long learned MAC addresses are stored. By default, each learned MAC address is aged out after 300 seconds. After this amount of time, the inactive MAC address is deleted from the FortiSwitch hardware. The value ranges from 10 to 1000,000 seconds. Set the value to 0 to disable MAC address aging.

```
config switch-controller global
  set mac-aging-interval <10 to 1000000>
end
```

For example:

```
config switch-controller global
  set mac-aging-interval 500
end
```

If you want to see the first MAC address that exceeded the learning limit for an interface or VLAN, you can enable the learning-limit violation log for a managed FortiSwitch unit. Only one violation is recorded per interface or VLAN.

By default, logging is disabled. The most recent violation that occurred on each interface or VLAN is recorded in the system log. After that, no more violations are logged until the log is reset for the triggered interface or VLAN. Only the most recent 128 violations are displayed in the console.

Use the following commands to control the learning-limit violation log and to control how long learned MAC addresses are save:

```
config switch-controller global
    set mac-violation-timer <0-1500>
    set log-mac-limit-violations {enable | disable}
end
```

For example:

```
config switch-controller global
    set mac-violation-timer 1000
    set log-mac-limit-violations enable
end
```

To view the content of the learning-limit violation log for a managed FortiSwitch unit, use one of the following commands:

- `diagnose switch-controller dump mac-limit-violations all <FortiSwitch_serial_number>`
- `diagnose switch-controller dump mac-limit-violations interface <FortiSwitch_serial_number> <port_name>`
- `diagnose switch-controller dump mac-limit-violations vlan <FortiSwitch_serial_number> <VLAN_ID>`

For example, to set the learning-limit violation log for VLAN 5 on a managed FortiSwitch unit:

```
diagnose switch-controller dump mac-limit-violations vlan S124DP3XS12345678 5
```

To reset the learning-limit violation log for a managed FortiSwitch unit, use one of the following commands:

- `execute switch-controller mac-limit-violation reset all <FortiSwitch_serial_number>`
- `execute switch-controller mac-limit-violation reset vlan <FortiSwitch_serial_number> <VLAN_ID>`
- `execute switch-controller mac-limit-violation reset interface <FortiSwitch_serial_number> <port_name>`

For example, to clear the learning-limit violation log for port 5 of a managed FortiSwitch unit:

```
execute switch-controller mac-limit-violation reset interface S124DP3XS12345678 port5
```

Configuring the DHCP trust setting

The DHCP blocking feature monitors the DHCP traffic from untrusted sources (for example, typically host ports and unknown DHCP servers) that might initiate traffic attacks or other hostile actions. To prevent this, DHCP blocking filters messages on untrusted ports.

Set the port as a trusted or untrusted DHCP-snooping interface:

```
config switch-controller managed-switch
  edit <switch-id>
    config ports
      edit <port name>
        set dhcp-snooping {trusted | untrusted}
      end
    end
  end
```

For example:

```
config switch-controller managed-switch
  edit S524DF4K15000024
    config ports
      edit port1
        set dhcp-snooping trusted
      end
    end
  end
```

Configuring PoE

The following PoE CLI commands are available starting in FortiSwitchOS 3.3.0.

Enable PoE on the port

```
config switch-controller managed-switch
  edit <switch-id>
    config ports
      edit <port name>
        set poe-status {enable | disable}
      end
    end
  end
```

For example:

```
config switch-controller managed-switch
  edit S524DF4K15000024
    config ports
      edit port1
        set poe-status enable
      end
    end
  end
```

Reset the PoE port

Power over Ethernet (PoE) describes any system that passes electric power along with data on twisted pair Ethernet cabling. Doing this allows a single cable to provide both data connection and electric power to devices (for example, wireless access points, IP cameras, and VoIP phones).

The following command resets PoE on the port:

```
execute switch-controller poe-reset <fortiswitch-id> <port>
```

Display general PoE status

```
get switch-controller <fortiswitch-id> <port>
```

The following example displays the PoE status for port 6 on the specified switch:

```
# get switch-controller poe FS108D3W14000967 port6
Port(6) Power:3.90W, Power-Status: Delivering Power
Power-Up Mode: Normal Mode
Remote Power Device Type: IEEE802.3AT PD
Power Class: 4
Defined Max Power: 30.0W, Priority:3
Voltage: 54.00V
Current: 78mA
```

Configuring edge ports

Use the following commands to enable or disable an interface as an edge port:

```
config switch-controller managed-switch
edit <switch>
config ports
edit <port>
set edge-port {enable | disable}
end
end
```

For example:

```
config switch-controller managed-switch
edit S524DF4K15000024
config ports
edit port1
set edge-port enable
end
end
```

Configuring STP

Starting with FortiSwitch Release 3.4.2, STP is enabled by default for the non-FortiLink ports on the managed FortiSwitch units. STP is a link-management protocol that ensures a loop-free layer-2 network topology.

NOTE: STP is not supported between a FortiGate unit and a FortiSwitch unit in FortiLink mode.

To configure global STP settings, see [Configure STP settings on page 1960](#).

Use the following commands to enable or disable STP on FortiSwitch ports:

```
config switch-controller managed-switch
edit <switch-id>
config ports
edit <port name>
set stp-state {enabled | disabled}
end
end
```

For example:

```
config switch-controller managed-switch
edit S524DF4K15000024
config ports
edit port1
set stp-state enabled
```

```

end
end

```

To check the STP configuration on a FortiSwitch, use the following command:

```
diagnose switch-controller dump stp <FortiSwitch_serial_number> <instance_number>
```

For example:

```
FG100D3G15817028 # diagnose switch-controller dump stp S524DF4K15000024 0
```

```
MST Instance Information, primary-Channel:
```

```
Instance ID : 0
```

```
Switch Priority : 24576
```

```
Root MAC Address : 085b0ef195e4
```

```
Root Priority: 24576
```

```
Root Pathcost: 0
```

```
Regional Root MAC Address : 085b0ef195e4
```

```
Regional Root Priority: 24576
```

```
Regional Root Path Cost: 0
```

```
Remaining Hops: 20
```

```
This Bridge MAC Address : 085b0ef195e4
```

```
This bridge is the root
```

Port Protection	Speed	Cost	Priority	Role	State	Edge	STP-Status	Loop
port1	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port2	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port3	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port4	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port5	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port6	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port7	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port8	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port9	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port10	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port11	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port12	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port13	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port14	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port15	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port16	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port17	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port18	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port19	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port20	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port21	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port22	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port23	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port25	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port26	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port27	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port28	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port29	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
port30	-	200000000	128	DISABLED	DISCARDING	YES	ENABLED	NO
internal	1G	20000	128	DESIGNATED	FORWARDING	YES	DISABLED	NO
__FoRtIlLiNk0__	1G	20000	128	DESIGNATED	FORWARDING	YES	DISABLED	NO

Configuring STP root guard

Root guard protects the interface on which it is enabled from becoming the path to root. When enabled on an interface, superior BPDUs received on that interface are ignored or dropped. Without using root guard, any switch

that participates in STP maintains the ability to reroute the path to root. Rerouting might cause your network to transmit large amounts of traffic across suboptimal links or allow a malicious or misconfigured device to pose a security risk by passing core traffic through an insecure device for packet capture or inspection. By enabling root guard on multiple interfaces, you can create a perimeter around your existing paths to root to enforce the specified network topology.

Enable root guard on all ports that should not be root bridges. Do not enable root guard on the root port. You must have STP enabled to be able to use root guard.

Use the following commands to enable or disable STP root guard on FortiSwitch ports:

```
config switch-controller managed-switch
  edit <switch-id>
    config ports
      edit <port name>
        set stp-root-guard {enabled | disabled}
      end
    end
  end
```

For example:

```
config switch-controller managed-switch
  edit S524DF4K15000024
    config ports
      edit port1
        set stp-root-guard enabled
      end
    end
  end
```

Configuring STP BPDU guard

Similar to root guard, BPDU guard protects the designed network topology. When BPDU guard is enabled on STP edge ports, any BPDUs received cause the ports to go down for a specified number of minutes. The BPDUs are not forwarded, and the network edge is enforced.

There are two prerequisites for using BPDU guard:

- You must define the port as an edge port with the `set edge-port enable` command.
- You must enable STP on the switch interface with the `set stp-state enabled` command.

You can set how long the port will go down when a BPDU is received for a maximum of 120 minutes. The default port timeout is 5 minutes. If you set the timeout value to 0, the port will not go down when a BPDU is received, but you will have manually reset the port.

Use the following commands to enable or disable STP BPDU guard on FortiSwitch ports:

```
config switch-controller managed-switch
  edit <switch-id>
    config ports
      edit <port name>
        set stp-bpdu-guard {enabled | disabled}
        set stp-bpdu-guard-time <0-120>
      end
    end
  end
```

For example:

```
config switch-controller managed-switch
```

```

edit S524DF4K15000024
  config ports
    edit port1
      set stp-bpdu-guard enabled
      set stp-bpdu-guard-time 10
    end
  end
end

```

To check the configuration of STP BPDU guard on a FortiSwitch unit, use the following command:

```
diagnose switch-controller dump bpdu-guard-status <FortiSwitch_serial_number>
```

For example:

```

FG100D3G15817028 # diagnose switch-controller dump bpdu-guard-status
S524DF4K15000024
Managed Switch : S524DF4K15000024 0

```

Portname	State	Status	Timeout (m)	Count	Last-Event
port1	enabled	-	10	0	-
port2	disabled	-	-	-	-
port3	disabled	-	-	-	-
port4	disabled	-	-	-	-
port5	disabled	-	-	-	-
port6	disabled	-	-	-	-
port7	disabled	-	-	-	-
port8	disabled	-	-	-	-
port9	disabled	-	-	-	-
port10	disabled	-	-	-	-
port11	disabled	-	-	-	-
port12	disabled	-	-	-	-
port13	disabled	-	-	-	-
port14	disabled	-	-	-	-
port15	disabled	-	-	-	-
port16	disabled	-	-	-	-
port17	disabled	-	-	-	-
port18	disabled	-	-	-	-
port19	disabled	-	-	-	-
port20	disabled	-	-	-	-
port21	disabled	-	-	-	-
port22	disabled	-	-	-	-
port23	disabled	-	-	-	-
port25	disabled	-	-	-	-
port26	disabled	-	-	-	-
port27	disabled	-	-	-	-
port28	disabled	-	-	-	-
port29	disabled	-	-	-	-
port30	disabled	-	-	-	-
__FortiLink0__	disabled	-	-	-	-

Configuring loop guard

A loop in a layer-2 network results in broadcast storms that have far-reaching and unwanted effects. Fortinet loop guard helps to prevent loops. When loop guard is enabled on a switch port, the port monitors its subtending network for any downstream loops. The loop guard feature is designed to work in concert with STP rather than as a replacement for STP. By default, loop guard is disabled on all ports.

Use the following commands to configure loop guard on a FortiSwitch port:

```
config switch-controller managed-switch
  edit <switch-id>
    config ports
      edit <port name>
        set loop-guard {enabled | disabled}
        set loop-guard-timeout <0-120 minutes>
      end
    end
  end
```

For example:

```
config switch-controller managed-switch
  edit S524DF4K15000024
    config ports
      edit port1
        set loop-guard enabled
        set loop-guard-timeout 10
      end
    end
  end
```

Configuring LLDP settings

The Fortinet data center switches support the Link Layer Discovery Protocol (LLDP) for transmission and reception wherein the switch will multicast LLDP packets to advertise its identity and capabilities. A switch receives the equivalent information from adjacent layer-2 peers.

Use the following commands to configure LLDP on a FortiSwitch port:

```
config switch-controller managed-switch
  edit <switch-id>
    config ports
      edit <port name>
        set lldp-status {rx-only | tx-only | tx-rx | disable}
        set lldp-profile <profile name>
      end
    end
  end
```

For example:

```
config switch-controller managed-switch
  edit S524DF4K15000024
    config ports
      edit port2
        set lldp-status tx-rx
        set lldp-profile default
      end
    end
  end
```

Configuring IGMP settings

IGMP snooping allows the FortiSwitch to passively listen to the Internet Group Management Protocol (IGMP) network traffic between hosts and routers. The switch uses this information to determine which ports are interested in receiving each multicast feed. FortiSwitch can reduce unnecessary multicast traffic on the LAN by pruning multicast traffic from links that do not contain a multicast listener.

Use the following commands to configure IGMP settings on a FortiSwitch port:

```
config switch-controller managed-switch
  edit <switch-id>
    config ports
      edit <port name>
        set igmp-snooping {enable | disable}
        set igmps-flood-reports {enable | disable}
      end
    end
  end
```

For example:

```
config switch-controller managed-switch
  edit S524DF4K15000024
    config ports
      edit port3
        set igmp-snooping enable
        set igmps-flood-reports enable
      end
    end
  end
```

Configuring sFlow

sFlow is a method of monitoring the traffic on your network to identify areas on the network that might impact performance and throughput. With sFlow, you can export truncated packets and interface counters. FortiSwitch implements sFlow version 5 and supports trunks and VLANs.

NOTE: Because sFlow is CPU intensive, Fortinet does not recommend high rates of sampling for long periods.

sFlow uses packet sampling to monitor network traffic. The sFlow agent captures packet information at defined intervals and sends them to an sFlow collector for analysis, providing real-time data analysis. To minimize the impact on network throughput, the information sent is only a sampling of the data.

The sFlow collector is a central server running software that analyzes and reports on network traffic. The sampled packets and counter information, referred to as flow samples and counter samples, respectively, are sent as sFlow datagrams to a collector. Upon receiving the datagrams, the sFlow collector provides real-time analysis and graphing to indicate the source of potential traffic issues. sFlow collector software is available from a number of third-party software vendors. You must configure a FortiGate policy to transmit the samples from the FortiSwitch unit to the sFlow collector.

sFlow can monitor network traffic in two ways:

- Flow samples—You specify the percentage of packets (one out of n packets) to randomly sample.
- Counter samples—You specify how often (in seconds) the network device sends interface counters.

Use the following CLI commands to specify the IP address and port for the sFlow collector. By default, the IP address is 0.0.0.0, and the port number is 6343.

```
config switch-controller sflow
  collector-ip <x.x.x.x>
```

```

    collector-port <port_number>
end

```

Use the following CLI commands to configure sFlow:

```

config switch-controller managed-switch <FortiSwitch_serial_number>
    config ports
        edit <port_name>
            set sflow-sampler <disabled | enabled>
            set sflow-sample-rate <0-99999>
            set sflow-counter-interval <1-255>
        next
    next
end

```

For example:

```

config switch-controller sflow
    collector-ip 1.2.3.4
    collector-port 10
end

config switch-controller managed-switch S524DF4K15000024
    config ports
        edit port5
            set sflow-sampler enabled
            set sflow-sample-rate 10
            set sflow-counter-interval 60
        next
    next
end

```

Configuring Dynamic ARP inspection (DAI)

DAI prevents man-in-the-middle attacks and IP address spoofing by checking that packets from untrusted ports have valid IP-MAC-address binding. DAI allows only valid ARP requests and responses to be forwarded.

To use DAI, you must first enable the DHCP-snooping feature, enable DAI, and then enable DAI for each VLAN. By default, DAI is disabled on all VLANs.

After enabling DHCP snooping with the `set switch-controller-dhcp-snooping enable` command, use the following CLI commands to enable DAI and then enable DAI for a VLAN:

```

config system interface
    edit vsw.test
        set switch-controller-arp-inspection <enable | disable>
    end

config switch-controller managed-switch
    edit <sn>
        config ports
            edit <VLAN_ID>
                arp-inspection-trust <untrusted | trusted>
            next
        end
    next
end

```

Use the following CLI command to check DAI statistics for a FortiSwitch unit:

```
diagnose switch arp-inspection stats <FortiSwitch_Serial_Number>
```

Use the following CLI command to delete DAI statistics for a specific VLAN:

```
diagnose switch arp-inspection stats clear <VLAN_ID> <FortiSwitch_Serial_Number>
```

Configuring FortiSwitch port mirroring

The FortiSwitch unit can send a copy of any ingress or egress packet on a port to egress on another port of the same FortiSwitch unit. The original traffic is unaffected. This process is known as port mirroring and is typically used for external analysis and capture.

Use the following CLI commands to configure FortiSwitch port mirroring:

```
config switch-controller managed-switch
  edit <FortiSwitch_Serial_Number>
    config mirror
      edit <mirror_name>
        set status <active | inactive>
        set dst <port_name>
        set switching-packet <enable | disable>
        set src-ingress <port_name>
        set src-egress <port_name>
      next
    end
  next
```

NOTE: The `set status` and `set dst` commands are mandatory for port mirroring.

For example:

```
config switch-controller managed-switch
  edit S524DF4K15000024
    config mirror
      edit 2
        set status active
        set dst port1
        set switching-packet enable
        set src-ingress port2 port3
        set src-egress port4 port5
      next
    end
  next
```

FortiSwitch port security policy

To control network access, the managed FortiSwitch unit supports IEEE 802.1x authentication. A supplicant connected to a port on the switch must be authenticated by a RADIUS/Diameter server to gain access to the network. The supplicant and the authentication server communicate using the switch using the EAP protocol. The managed FortiSwitch unit supports EAP-PEAP, EAP-TTLS, EAP-TLS, and EAP-MD5.

To use the RADIUS server for authentication, you must configure the server before configuring the users or user groups on the managed FortiSwitch unit.

NOTE: In FortiLink mode, you must manually create a firewall policy to allow RADIUS traffic for 802.1x authentication from the FortiSwitch unit (for example, from the FortiLink interface) to the RADIUS server through the FortiGate.

The managed FortiSwitch unit implements MAC-based authentication. The switch saves the MAC address of each supplicant's device. The switch provides network access only to devices that have successfully been authenticated.

You can enable the MAC Authentication Bypass (MAB) option for devices (such as network printers) that cannot respond to the 802.1x authentication request. With MAB enabled on the port, the system will use the device MAC address as the user name and password for authentication.

Optionally, you can configure a guest VLAN for unauthorized users. Alternatively, you can specify a VLAN for users whose authentication was unsuccessful.

When you are testing your system configuration for 802.1x authentication, you can use the monitor mode to allow network traffic to flow, even if there are configuration problems or authentication failures.

This chapter covers the following topics:

- ["Configure the 802.1X settings for a virtual domain" on page 1982](#)
- ["Override the virtual domain settings" on page 1983](#)
- ["Define an 802.1X security policy" on page 1983](#)
- ["Apply an 802.1X security policy to a FortiSwitch port" on page 1985](#)
- ["Test 802.1x authentication with monitor mode" on page 1985](#)
- ["Restrict the type of frames allowed through IEEE 802.1Q ports" on page 1986](#)
- ["RADIUS accounting support" on page 1986](#)

Configure the 802.1X settings for a virtual domain

To configure the 802.1X security policy for a virtual domain, use the following commands:

```
config switch-controller 802-1X-settings
    set reauth-period < int >
    set max-reauth-attempt < int >
    set link-down-auth < *set-unauth | no-action >
end
```

Option	Description
<code>set link-down-auth</code>	If a link is down, this command determines the authentication state. Choosing <code>set-auth</code> sets the interface to unauthenticated when a link is down, and reauthentication is needed. Choosing <code>no-auth</code> means that the interface does not need to be reauthenticated when a link is down.
<code>set reauth-period</code>	This command sets how often reauthentication is needed. The range is 1-1440 minutes. The default is 60 minutes. Setting the value to 0 minutes disables reauthentication.

Option	Description
<code>set max-reauth-attempt</code>	This command sets the maximum number of reauthentication attempts. The range is 1-15. the default is 3. Setting the value to 0 disables reauthentication.

Override the virtual domain settings

You can override the virtual domain settings for the 802.1X security policy.

Using the FortiGate GUI

To override the 802.1X settings for a virtual domain:

1. Go to *WiFi & Switch Controller > Managed FortiSwitch*.
2. Click on a FortiSwitch faceplate and select *Edit*.
3. In the Edit Managed FortiSwitch page, move the *Override 802-1X settings* slider to the right.
4. In the Reauthentication Interval field, enter the number of minutes before reauthentication is required. The maximum interval is 1,440 minutes. Setting the value to 0 minutes disables reauthentication.
5. In the Max Reauthentication Attempts field, enter the maximum times that reauthentication is attempted. The maximum number of attempts is 15. Setting the value to 0 disables reauthentication.
6. Select *Deauthenticate* or *None* for the link down action. Selecting *Deauthenticate* sets the interface to unauthenticated when a link is down, and reauthentication is needed. Selecting *None* means that the interface does not need to be reauthenticated when a link is down.
7. Select *OK*.

Using the FortiGate CLI

To override the 802.1X settings for a virtual domain, use the following commands:

```
config switch-controller managed-switch
  edit < switch >
    config 802-1X-settings
      set local-override [ enable | *disable ]
      set reauth-period < int >                // visible if override enabled
      set max-reauth-attempt < int >           // visible if override enabled
      set link-down-auth < *set-unauth | no-action > // visible if override enabled
    end
  next
end
```

For a description of the options, see [Configure the 802.1X settings for a virtual domain](#).

Define an 802.1X security policy

You can define multiple 802.1X security policies.

Using the FortiGate GUI

To create an 802.1X security policy:

1. Go to *WiFi & Switch Controller > FortiSwitch Security Policies*.
2. Select *Create New*.
3. Enter a name for the new FortiSwitch security policy.
4. For the security mode, select *Port-based* or *MAC-based*.
5. Select + to select which user groups will have access.
6. Enable or disable guest VLANs on this interface to allow restricted access for some users.
7. Enter the number of seconds for authentication delay for guest VLANs. The range is 1-900 seconds.
8. Enable or disable authentication fail VLAN on this interface to allow restricted access for users who fail to access the guest VLAN.
9. Enable or disable MAC authentication bypass (MAB) on this interface.
10. Enable or disable EAP pass-through mode on this interface.
11. Enable or disable whether the session timeout for the RADIUS server will overwrite the local timeout.
12. Select *OK*.

Using the FortiGate CLI

To create an 802.1X security policy, use the following commands:

```
config switch-controller security-policy 802-1X
edit "<policy.name>"
    set security-mode {802.1X | 802.1X-mac-based}
    set user-group <*group_name | Guest-group | SSO_Guest_Users>
    set mac-auth-bypass [enable | *disable]
    set eap-passthru [enable | disable]
    set guest-vlan [enable | *disable]
    set guest-vlan-id "guest-VLAN-name"
    set guest-auth-delay <integer>
    set auth-fail-vlan [enable | *disable]
    set auth-fail-vlan-id "auth-fail-VLAN-name"
    set radius-timeout-overwrite [enable | *disable]
    set policy-type 802.1X
end
end
```

Option	Description
<code>set security-mode</code>	You can restrict access with 802.1X port-based authentication or with 802.1X MAC-based authentication.
<code>set user-group</code>	You can set a specific group name, Guest-group, or SSO_Guest_Users to have access. This setting is mandatory.
<code>set mac-auth-bypass</code>	You can enable or disable MAB on this interface.
<code>set eap-passthrough</code>	You can enable or disable EAP pass-through mode on this interface.
<code>set guest-vlan</code>	You can enable or disable guest VLANs on this interface to allow restricted access for some users.

Option	Description
<code>set guest-vlan-id "guest-VLAN-name"</code>	You can specify the name of the guest VLAN.
<code>set guest-auth-delay</code>	You can set the authentication delay for guest VLANs on this interface. The range is 1-900 seconds.
<code>set auth-fail-vlan</code>	You can enable or disable authentication fail VLAN on this interface to allow restricted access for users who fail to access the guest VLAN.
<code>set auth-fail-vlan-id "auth-fail-VLAN-name"</code>	You can specify the name of the authentication fail VLAN
<code>set radius-timeout-overwrite</code>	You can enable or disable whether the session timeout for the RADIUS server will overwrite the local timeout.
<code>set policy-type 802.1X</code>	You can set the policy type to the 802.1X security policy.

Apply an 802.1X security policy to a FortiSwitch port

You can apply a different 802.1X security policy to each FortiSwitch port.

Using the FortiGate GUI

To apply an 802.1X security policy to a managed FortiSwitch port:

1. Go to *WiFi & Switch Controller > FortiSwitch Ports*.
2. Select the + next to a FortiSwitch unit.
3. In the Security Policy column for a port, click + to select a security policy.
4. Select *OK* to apply the security policy to that port.

Using the FortiGate CLI

To apply an 802.1X security policy to a managed FortiSwitch port, use the following commands:

```
config switch-controller managed-switch
  edit <managed-switch>
    config ports
      edit <port>
        set port-security-policy <802.1X-policy>
      next
    end
  next
end
```

Test 802.1x authentication with monitor mode

Use the monitor mode to test your system configuration for 802.1x authentication. You can use monitor mode to test port-based authentication, MAC-based authentication, EAP pass-through mode, and MAC authentication

bypass. Monitor mode is disabled by default. After you enable monitor mode, the network traffic will continue to flow, even if the users fail authentication.

To enable or disable monitor mode, use the following commands:

```
config switch-controller security-policy 802-1X
  edit "<policy_name>"
    set open-auth {enable | disable}
  next
end
```

Restrict the type of frames allowed through IEEE 802.1Q ports

You can now specify whether each FortiSwitch port discards tagged 802.1Q frames or untagged 802.1Q frames or allows all frames access to the port. By default, all frames have access to each FortiSwitch port.

Use the following CLI commands:

```
config switch-controller managed-switch <SN>
  config ports
    edit <port_name>
      set discard-mode <none | all-tagged | all-untagged>
    next
  next
end
```

RADIUS accounting support

The FortiSwitch unit uses 802.1x-authenticated ports to send five types of RADIUS accounting messages to the RADIUS accounting server to support FortiGate RADIUS single sign-on:

- **START**—The FortiSwitch has been successfully authenticated, and the session has started.
- **STOP**—The FortiSwitch session has ended.
- **INTERIM**—Periodic messages sent based on the value set using the `set acct-interim-interval` command.
- **ON**—FortiSwitch will send this message when the switch is turned on.
- **OFF**—FortiSwitch will send this message when the switch is shut down.

Use the following commands to set up RADIUS accounting so that FortiOS can send accounting messages to managed FortiSwitch units:

```
config user radius
  edit <RADIUS_server_name>
    set acct-interim-interval <seconds>
    config accounting-server
      edit <entry_ID>
        set status {enable | disable}
        set server <server_IP_address>
        set secret <secret_key>
        set port <port_number>
      next
    end
  next
end
```

Additional capabilities

This chapter covers the following topics:

- [Execute custom FortiSwitch commands on page 1987](#)
- [View and upgrade the FortiSwitch firmware version on page 1988](#)
- [FortiSwitch log export on page 1989](#)
- [FortiSwitch per-port device visibility on page 1989](#)
- [FortiGate CLI support for FortiSwitch features \(on non-FortiLink ports\) on page 1989](#)
- [Synchronizing the FortiGate unit with the managed FortiSwitch units on page 1994](#)
- [Replacing a managed FortiSwitch unit on page 1995](#)

Execute custom FortiSwitch commands

From the FortiGate, you can execute FortiSwitch commands on the managed FortiSwitch.

This feature adds a simple scripting mechanism for users to execute generic commands on the switch.

NOTE: FortiOS 5.6.0 introduces additional capabilities related to the managed FortiSwitch.

Create a command

Use the following syntax to create a command file:

```
config switch-controller custom-command
  edit <cmd-name>
    set command " <FortiSwitch commands>"
```

Next, create a command file to set the STP max-age parameter:

```
config switch-controller custom-command
  edit "stp-age-10"
    set command "config switch stp setting
      set max-age 10
    end
  "
next
end
```

Execute a command

After you have created a command file, use the following command on the FortiGate to execute the command file on the target switch:

```
exec switch-controller custom-command <cmd-name> <target-switch>
```

The following example runs the **stp-age-10** command on the specified target FortiSwitch:

```
# exec switch-controller custom-command stp-age-10 S124DP3X15000118
```

View and upgrade the FortiSwitch firmware version

You can view the current firmware version of a FortiSwitch and upgrade the FortiSwitch to a new firmware version. FortiGate will suggest an upgrade when a new version is available in FortiGuard.

Using the FortiGate Web interface

To view the FortiSwitch firmware version:

1. Go to **WiFi & Switch Controller>Managed FortiSwitch**.
2. In the main panel, select the FortiSwitch faceplate and click **Edit**.
3. In the **Edit Managed FortiSwitch** panel, the **Firmware** section displays the current build on the FortiSwitch.

To upgrade the firmware on multiple FortiSwitch units at the same time:

1. Go to *WiFi & Switch Controller > Managed FortiSwitch*.
2. Select the faceplates of the FortiSwitch units that you want to upgrade.
3. Click *Upgrade*.
The *Upgrade FortiSwitches* page opens.
4. Select *FortiGuard* or select *Upload* and then select the firmware file to upload.
If you select *FortiGuard*, all FortiSwitch units that can be upgraded are upgraded. If you select *Upload*, only one firmware image can be used at a time for upgrading.
5. Select *Upgrade*.

Using the CLI

Use the following command to display the latest version:

```
diagnose fdsm fortisw-latest-ver <model>
```

Use the following command to download the image:

```
diagnose fdsm fortisw-download <image id>
```

The following example shows how to download the latest image for FS224D:

```
FG100D3G15801204 (global) # diagnose fdsm fortisw-latest-ver FS224D
FS224D - 3.4.2 b192 03004000FIMG0900904002FG100D3G15801204 (global) #

diagnose fdsm fortisw-download 03004000FIMG0900904002

Download image-03004000FIMG0900904002:
#####
Result=Success
```

Use the following CLI commands to enable the use of HTTPS to download firmware to managed FortiSwitch units:

```
config switch-controller global
  set https-image-push enable
end
```

From your FortiGate CLI, you can upgrade the firmware of all of the managed FortiSwitch units of the same model using a single `execute` command. The command includes the name of a firmware image file and all of the managed FortiSwitch units compatible with that firmware image file are upgraded. For example:

```
execute switch-controller stage-tiered-swtp-image ALL <firmware-image-file>
```

You can also use the following command to restart all of the managed FortiSwitch units after a 2-minute delay.

```
execute switch-controller restart-swtp-delayed ALL
```

FortiSwitch log export

You can enable and disable the managed FortiSwitch units to export their syslogs to the FortiGate. The setting is global, and the default setting is enabled. Starting in FortiOS 5.6.3, more details are included in the exported FortiSwitch logs.

To allow a level of filtering, FortiGate sets the user field to "fortiswitch-syslog" for each entry.

The following is the CLI command syntax:

```
config switch-controller switch-log
  set status (*enable | disable)
  set severity [emergency | alert | critical | error | warning | notification |
    *information | debug]
end
```

You can override the global log settings for a FortiSwitch, using the following commands:

```
config switch-controller managed-switch
  edit <switch-id>
    config switch-log
      set local-override enable
```

At this point, you can configure the log settings that apply to this specific switch.

FortiSwitch per-port device visibility

In the FortiGate GUI, **User & Device > Device List** displays a list of devices attached to the FortiSwitch ports. For each device, the table displays the IP address of the device and the interface (FortiSwitch name and port).

From the CLI, the following command displays information about the host devices:

```
diagnose switch-controller dump mac-hosts_switch-ports
```

FortiGate CLI support for FortiSwitch features (on non-FortiLink ports)

You can configure the following FortiSwitch features from the FortiGate CLI.

Configuring a link aggregation group (LAG)

You can configure a link aggregation group (LAG) for non-FortiLink ports on a FortiSwitch. You cannot configure ports from different FortiSwitch units in one LAG.

```
config switch-controller managed-switch
  edit <switch-id>
    config ports
      it <trunk name>
```

```

        set type trunk
        set mode < static | lacp > Link Aggregation mode
        set bundle (enable | disable)
        set min-bundle <int>
        set max-bundle <int>
        set members < port1 port2 ...>
    next
end
end
end
end

```

Configuring an MCLAG with managed FortiSwitch units

A multichassis LAG (MCLAG) provides node-level redundancy by grouping two FortiSwitch models together so that they appear as a single switch on the network. If either switch fails, the MCLAG continues to function without any interruption, increasing network resiliency and eliminating the delays associated with the Spanning Tree Protocol (STP). For the network topology, see [Dual-homed servers connected to FortiLink tier-1 FortiSwitch units using an MCLAG on page 1938](#) and [Standalone FortiGate unit with dual-homed FortiSwitch access on page 1939](#).

Notes

- Both peer switches should be of the same hardware model and same software version. Mismatched configurations might work but are unsupported.
- There is a maximum of two FortiSwitch models per MCLAG.
- The routing feature is not available within an MCLAG.
- For static MAC addresses within an MCLAG, if one FortiSwitch learns the MAC address, the second FortiSwitch will automatically learn the MAC address.

To configure an MCLAG with managed FortiSwitch units:

1. For each MCLAG peer switch, log into the FortiSwitch to create a LAG:

```

config switch trunk
    edit "LAG-member"
        set mode lacp-active
        set mclag-icl enable
        set members "<port>" "<port>"
    next

```

2. Enable the MCLAG on each managed FortiSwitch:

```
config switch-controller managed-switch
  edit "<switch-id>"
    config ports
      edit "<trunk name>"
        set type trunk
        set mode {static | lacp-passive | lacp-active}
        set bundle {enable | disable}
        set members "<port>,<port>"
        set mclag {enable | disable}
      next
    end
  next
```

3. Log into each managed FortiSwitch to check the MCLAG configuration:

```
diagnose switch mclag
```

After the FortiSwitch units are configured as MCLAG peer switches, any port that supports advanced features on the FortiSwitch can become a LAG port. When `mclag` is enabled and the LAG port names match, an MCLAG peer set is automatically formed. The member ports for each FortiSwitch in the MCLAG do not need to be identical to the member ports on the peer FortiSwitch.

Configuring storm control

Storm control uses the data rate (packets/sec, default 500) of the link to measure traffic activity, preventing traffic on a LAN from being disrupted by a broadcast, multicast, or unicast storm on a port.

When the data rate exceeds the configured threshold, storm control drops excess traffic. You can configure the types of traffic to drop: broadcast, unknown unicast, or multicast.

The storm control settings are global to all of the non-FortiLink ports on the managed switches. Use the following CLI commands to configure storm control:

```
config switch-controller storm-control
  set rate <rate>
  set unknown-unicast {enable | disable}
  set unknown-multicast {enable | disable}
  set broadcast {enable | disable}
end
```

You can override the global storm control settings for a FortiSwitch using the following commands:

```
config switch-controller managed-switch
  edit <switch-id>
    config storm-control
      set local-override enable
```

At this point, you can configure the storm control settings that apply to this specific switch.

Displaying port statistics

Port statistics will be accessed using the following FortiSwitch CLI command:

```
FG100D3G15804763 # diagnose switch-controller dump port-stats
S124DP3X16000413 port8
```

```

S124DP3X16000413 0 :
{
  "port8":{
    "tx-bytes":823526672,
    "tx-packets":1402390,
    "tx-ucast":49047,
    "tx-mcast":804545,
    "tx-bcast":548798,
    "tx-errors":0,
    "tx-drops":3,
    "tx-oversize":0,
    "rx-bytes":13941793,
    "rx-packets":160303,
    "rx-ucast":148652,
    "rx-mcast":7509,
    "rx-bcast":4142,
    "rx-errors":0,
    "rx-drops":720,
    "rx-oversize":0,
    "undersize":0,
    "fragments":0,
    "jabbers":0,
    "collisions":0,
    "crc-alignments":0,
    "l3packets":0
  }
}

```

Configuring QoS with managed FortiSwitch units

Quality of Service (QoS) provides the ability to set particular priorities for different applications, users, or data flows.

NOTE: The FortiGate unit does not support QoS for hard or soft switch ports.

The FortiSwitch unit supports the following QoS configuration capabilities:

- Mapping the IEEE 802.1p and Layer 3 QoS values (Differentiated Services and IP Precedence) to an outbound QoS queue number.
- Providing eight egress queues on each port.
- Policing the maximum data rate of egress traffic on the interface.

To configure the QoS for managed FortiSwitch units:

1. Configure a Dot1p map.

A Dot1p map defines a mapping between IEEE 802.1p class of service (CoS) values (from incoming packets on a trusted interface) and the egress queue values. Values that are not explicitly included in the map will follow the default mapping, which maps each priority (0-7) to queue 0. If an incoming packet contains no CoS value, the switch assigns a CoS value of zero.

NOTE: Do not enable trust for both Dot1p and DSCP at the same time on the same interface. If you do want to trust both Dot1p and IP-DSCP, the FortiSwitch uses the latter value (DSCP) to determine the queue. The switch will use the Dot1p value and mapping only if the packet contains no DSCP value.

```
config switch-controller qos dot1p-map
```

```
edit <Dot1p map name>
  set description <text>
  set priority-0 <queue number>
  set priority-1 <queue number>
  set priority-2 <queue number>
  set priority-3 <queue number>
  set priority-4 <queue number>
  set priority-5 <queue number>
  set priority-6 <queue number>
  set priority-7 <queue number>
next
end
```

2. Configure a DSCP map.

A DSCP map defines a mapping between IP precedence or DSCP values and the egress queue values. For IP precedence, you have the following choices:

- network-control—Network control
- internetwork-control—Internetwork control
- critic-ecp—Critic and emergency call processing (ECP)
- flashoverride—Flash override
- flash—Flash
- immediate—Immediate
- priority—Priority
- routine—Routine

```
config switch-controller qos ip-dscp-map
edit <DSCP map name>
  set description <text>
  configure map <map_name>
  edit <entry name>
    set cos-queue <COS queue number>
    set diffserv {CS0 | CS1 | AF11 | AF12 | AF13 | CS2 | AF21 | AF22 | AF23
      | CS3 | AF31 | AF32 | AF33 | CS4 | AF41 | AF42 | AF43 | CS5 | EF |
      CS6 | CS7}
    set ip-precedence {network-control | internetwork-control | critic-ecp
      | flashoverride | flash | immediate | priority | routine}
    set value <DSCP raw value>
  next
end
end
```

3. Configure the egress QoS policy.

In a QoS policy, you set the scheduling mode for the policy and configure one or more CoS queues. Each egress port supports eight queues, and three scheduling modes are available:

- With strict scheduling, the queues are served in descending order (of queue number), so higher number queues receive higher priority.
- In simple round-robin mode, the scheduler visits each backlogged queue, servicing a single packet from each queue before moving on to the next one.
- In weighted round-robin mode, each of the eight egress queues is assigned a weight value ranging from 0 to 63.


```

config switch-controller qos queue-policy
  edit <QoS egress policy name>
    set schedule {strict | round-robin | weighted}
    config cos-queue
      edit [queue-<number>]
        set description <text>
        set min-rate <rate in kbps>
        set max-rate <rate in kbps>
        set drop-policy {taildrop | random-early-detection}
        set weight <weight value>
      next
    end
  next
end

```

4. Configure the overall policy that will be applied to the switch ports.

```

config switch-controller qos qos-policy
  edit <QoS egress policy name>
    set default-cos <default CoS value 0-7>
    set trust-dot1p-map <Dot1p map name>
    set trust-ip-dscp-map <DSCP map name>
    set queue-policy <queue policy name>
  next
end

```

5. Configure each switch port.

```

config switch-controller managed-switch
  edit <switch-id>
    config ports
      edit <port>
        set qos-policy <CoS policy>
      next
    end
  next
end

```

Synchronizing the FortiGate unit with the managed FortiSwitch units

You can synchronize the FortiGate unit with the managed FortiSwitch units to check for synchronization errors on each managed FortiSwitch unit.

Use the following command to synchronize the full configuration of a FortiGate unit with the managed FortiSwitch unit:

```
execute switch-controller trigger-config-sync <FortiSwitch_serial_number>
```

Use one of the following commands to display the synchronization state of a FortiGate unit with a specific managed FortiSwitch unit:

```

execute switch-controller get-sync-status switch-id <FortiSwitch_serial_number>
execute switch-controller get-sync-status name <FortiSwitch_name>

```

Use the following command to display the synchronization state of a FortiGate unit with a group of managed FortiSwitch units:

```
execute switch-controller get-sync-status group <FortiSwitch_group_name>
```

Use the following command to check the synchronization state of all managed FortiSwitch units in the current VDOM:

```
execute switch-controller get-sync-status all
```

For example:

```
FG100D3G14813513 (root) # execute switch-controller get-sync-status all
Managed-devices in current vdom root:
```

STACK-NAME: FortiSwitch-Stack-port5				
SWITCH (NAME)	STATUS	CONFIG	MAC-SYNC	UPGRADE
FS1D243Z14000173	Up	Idle	Idle	Idle
S124DP3X16006228 (Desktop-Switch)	Up	Idle	Idle	Idle

Replacing a managed FortiSwitch unit

If a managed FortiSwitch unit has a hardware issue, you can replace it with another FortiSwitch unit with the same model. The replacement FortiSwitch unit will inherit the configuration of the FortiSwitch unit that it replaces.

To replace a managed FortiSwitch unit:

1. Go to *WiFi & Switch Controller > Managed FortiSwitch*.
2. Select the faceplate of the FortiSwitch unit with the hardware issue.
3. Select *Deauthorize*.
4. In the CLI, enter the following command:

```
execute replace-device fortiswitch <existing_FortiSwitch_serial_number> <replacement_
FortiSwitch_serial_number>
```

For example:

```
execute replace-device fortiswitch S124DN3W16002025 S124DN3W16002026
```

NOTE: The two FortiSwitch serial numbers must belong to the same model.

Troubleshooting

Troubleshooting FortiLink issues

If the FortiGate does not establish the FortiLink connection with the FortiSwitch, perform the following troubleshooting checks.

Check the FortiGate configuration

To use the FortiGate GUI to check the FortiLink interface configuration:

1. In *Network > Interfaces*, double-click the interface used for FortiLink.
2. Ensure that *Dedicated to FortiSwitch* is set for this interface.

To use the FortiGate CLI to verify that you have configured the DHCP and NTP settings correctly:

1. Verify that the NTP server is enabled and that the FortiLink interface has been added to the list:

```
show system ntp
```

2. Ensure that the DHCP server on the Fortilink interface is configured correctly:

```
show system dhcp
```

Check the FortiSwitch configuration

To use FortiSwitch CLI commands to check the FortiSwitch configuration:

1. Verify that the switch system time matches the time on the FortiGate:

```
get system status
```

2. Verify that FortiGate has sent an IP address to the FortiSwitch (anticipate an IP address in the range 169.254.x.x):

```
get system interfaces
```

3. Verify that you can ping the FortiGate IP address:

```
exec ping x.x.x.x
```

To use FortiGate CLI commands to check the FortiSwitch configuration:

1. Verify that the connections from the FortiGate to the FortiSwitch units are up:

```
exec switch-controller get-conn-status
```

2. Verify that ports for a specific FortiSwitch stack are connected to the correct locations:

```
exec switch-controller get-physical-conn <FortiSwitch-Stack-ID>
```

3. Verify that all the ports for a specific FortiSwitch are up:

```
exec switch-controller get-conn-status <FortiSwitch-device-ID>
```

Check FortiSwitch connections

Use the following CLI command for detailed diagnostic information on the managed FortiSwitch connections:

```
execute switch-controller diagnose-connection <FortiSwitch_serial_number>
```

If the FortiSwitch serial number is omitted, only the FortiLink configuration is checked.

Chapter 17 - Networking

Introduction

This document explains how to configure your network. The following chapters are included in this document:

[Interfaces](#) explains the concepts of options for setting up interfaces and groupings of subnetworks that can scale to a company's growing requirements.

[DNS](#) explains how to set DNS requirements for your network and how to set FortiGate as a local DNS server.

[Advanced static routing](#) explains universal and static routing concepts, equal cost multipath (ECMP) and load balancing, policy routing, and routing in transparent mode.

[Dynamic routing](#) provides an overview of dynamic routing, compares static and dynamic routing, and describes dynamic routing protocols (RIP, OSPF, BGP, and IS-IS).

[Multicast forwarding](#) explains the concepts and use of multicasting with FortiGate.

[SD-WAN](#) describes FortiOS 6.0 features for SD-WAN.

[Troubleshooting](#) describes features, such as packet capture, that are useful for troubleshooting purposes.

What's new in FortiOS 6.0

The following list contains new Networking features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- ["Range for DHCPv6 prefix delegation" on page 2006](#)
- ["Configuring TFTP servers" on page 2008](#)
- TWAMP enhancements, see ["TWAMP" on page 2024](#)
- ["Link health monitor" on page 2025](#)
- ["VLANs over VXLANs" on page 2043](#)
- ["Enhanced MAC VLANs" on page 2049](#)
- ["VLAN filters for virtual wire pairs" on page 2053](#)
- Internet services enhancements, see ["Internet services" on page 2061](#)
- ["Static routes and VRFs" on page 2073](#)
- ["Source prefixes for static routes in transparent mode" on page 2084](#)
- IPv6 BFD support, see ["BFD" on page 2106](#)
- ["BFD and static routes" on page 2107](#)
- ["OSPF and VRFs" on page 2158](#)
- IS-IS support for IPv6, see ["IS-IS" on page 2224](#)
- Support for loopback interfaces with IS-IS, see ["Loopback interfaces" on page 2229](#)
- ["SD-WAN" on page 2273](#)

Interfaces

Interfaces, both physical and virtual, allow traffic to flow between internal networks and the Internet, and between internal networks. FortiGate has a number of options for setting up interfaces and groupings of subnetworks that can scale to your organization's growing requirements.

Administrative access

To help prevent FortiGate interfaces, especially the public-facing ports, from being accessed by users who you don't want accessing them, you can configure protocols that an administrator must use to access FortiGate, including:

- HTTPS
- PING
- FortiManager Access (FMG-Access)
- CAPWAP
- SSH
- SNMP
- FTM
- RADIUS Accounting
- FortiTelemetry

As a best practice, you should configure administrative access when you're setting the IP address for the port.

The following example adds the IPv4 address 172.20.120.100 to the WAN1 interface, and administrative access to HTTPS and SSH.

Add an IP address to the WAN1 interface - GUI

1. Go to **Network > Interfaces**.
2. Select the WAN1 interface row and select **Edit**.
3. Select the **Addressing Mode** of **Manual**.
4. Enter the IP address for the port of 172.20.120.100/24.
5. For **Administrative Access**, select **HTTPS** and **SSH**.
6. Select **OK**.

Add an IP address to the WAN1 interface - CLI

```
config system interface
  edit wan1
    set ip 172.20.120.100/24
    set allowaccess https ssh
  next
end
```



When you add or remove a protocol, you must type the entire list of protocols again. For example, if you have an access list of HTTPS and SSH and you want to add PING, you must use the following CLI command:

```
set allowaccess https ssh ping
```

If you use `set allowsaccess ping`, only ping is set and HTTPS and SSH are removed.

Aggregate interfaces

Link aggregation (IEEE 802.3ad) allows you to bind two or more physical interfaces together to form an aggregated link. This new link has the bandwidth of all the links combined. If a link in the group fails, traffic is automatically transferred to the remaining interfaces with the only noticeable effect being reduced bandwidth.

This is similar to redundant interfaces, with the major difference being that a redundant interface group uses only one link at a time, while an aggregate link group uses the total bandwidth of the functioning links in the group, up to eight (or more).

Some FortiGate models support the IEEE standard 802.3ad for link aggregation.

An interface can be an aggregate interface if it meets the following criteria:

- It's a physical interface, not a VLAN interface or subinterface.
- It's not already part of an aggregate or redundant interface.
- It's in the same VDOM as the aggregated interface. Aggregate ports can't span multiple VDOMs.
- It doesn't have an IP address and isn't configured for DHCP or PPPoE.
- It's not referenced in any security policy, VIP, IP pool, or multicast policy.
- It's not an HA heartbeat interface.
- It's not one of the backplane interfaces of the FortiGate 5000 series.

Some FortiGate models don't support aggregate interfaces. In this case, the aggregate option isn't available in the FortiGate GUI or CLI. Also, you can't create aggregate interfaces from interfaces in a switch port.

To see if a port is being used or has other dependencies - CLI:

```
diagnose sys checkused system.interface.name <interface_name>
```

When an interface is included in an aggregate interface, it's not listed in the **Network > Interfaces** page in the FortiGate GUI. Interfaces still appear in the CLI, but if you configure those interfaces, it won't take effect. You can't configure the interface individually and it's not available to include in security policies, VIPs, IP pools, or routing.

To avoid unintentional network issues when you configure Link Aggregation Control Protocol (LACP), disconnect the interfaces that you want to add to the aggregate interface. After you finish configuring LACP, reconnect the interfaces.

The following example creates an aggregate interface on a FortiGate, using ports 4 to 6, with an internal IP address of 10.13.101.100, and administrative access to HTTPS and SSH.

To create an aggregate interface - GUI:

1. Go to **Network > Interfaces** and select **Create New**, then **Interface**.
2. Enter the Name as `Aggregate`.
3. For the **Type**, select **802.3ad Aggregate**.
If this option doesn't appear, the FortiGate doesn't support aggregate interfaces.
4. In the **Physical Interface Members**, click to add interfaces. Select port 4, 5, and 6.
5. Select the **Addressing Mode** of **Manual**.
6. Enter the IP address for the port of 10.13.101.100/24.
7. For **Administrative Access**, select HTTPS and SSH.
8. Select **OK**.

To create aggregate interface - CLI:

```
config system interface
edit aggregate
set type aggregate
set member port4 port5 port6
set vdom root
set ip 172.20.120.100/24
set allowaccess https ssh
next
end
```

Sending GARP on aggregate MAC changes

A FortiGate sends out Gratuitous Address Resolution Protocol (GARP) announcements if the MAC address of a link aggregated interface changes to a new IP pool address due to a link failure or change in ports. This is needed when you use networking devices, such as some switches that don't perform this function when they receive LACP (Link Aggregation Control Protocol) information about changes in the MAC information.

DHCP addressing mode on an interface

If you configure an interface to use DHCP, FortiGate automatically broadcasts a DHCP request from the interface. The interface is configured with the IP address, any DNS server addresses, and the default gateway address that the DHCP server provides.

DHCP IPv6 is similar to DHCP IPv4, except:



- No default gateway option is defined because a host learns the gateway using router advertisement messages.
- There are no WINS servers because it is obsolete.

For more information about DHCP IPv6, see [RFC 3315](#).

You can configure DHCP for an interface in **Network > Interfaces** in the FortiGate GUI. Select the interface from the list, and select **DHCP** in the **Addressing mode**. The following table describes the DHCP status information when DHCP is configured for an interface.

Field	Description
Status	<p>Displays DHCP status messages as the interface connects to the DHCP server and gets addressing information. Select Status to refresh the addressing mode status message.</p> <p>Status can be one of the following values:</p> <ul style="list-style-type: none"> • initializing: no activity • connecting: interface attempts to connect to the DHCP server • connected: interface retrieves an IP address, netmask, and other settings from the DHCP server • failed: interface was unable to retrieve an IP address and other settings from the DHCP server
Obtained IP/Netmask	The IP address and netmask leased from the DHCP server. This is only displayed if the Status is connected .
Renew	Select this to renew the DHCP license for this interface. This is only displayed if the Status is connected .
Expiry Date	The time and date when the leased IP address and netmask is no longer valid for the interface. The IP address is returned to the pool to be allocated to the next user request for an IP address. This is only displayed if the Status is connected .
Default Gateway	The IP address of the gateway defined by the DHCP server. This is displayed only if the Status is connected , and if Receive default gateway from server is selected.
Distance	Enter the administrative distance for the default gateway retrieved from the DHCP server. The administrative distance is an integer from 1 to 255, and specifies the relative priority of a route when there are multiple routes to the same destination. A lower administrative distance indicates a more preferred route.
Retrieve default gateway from server	Enable this to retrieve a default gateway IP address from the DHCP server. The default gateway is added to the static routing table.
Override internal DNS	<p>Enable this to use the DNS addresses retrieved from the DHCP server instead of the DNS server IP addresses on the DNS page.</p> <p>When VDOMs are enabled, you can override the internal DNS only on the management VDOM.</p>

DHCP servers and relays

A DHCP server provides an address, from a defined address range, to a client on the network that requests it.

An interface can't provide both a server and a relay for connections of the same type (regular or IPsec). However, you can configure a regular DHCP server on an interface only if the interface is a physical interface with a static IP address. You can configure an IPsec DHCP server on an interface that has either a static or a dynamic IP address.

You can configure one or more DHCP servers on any FortiGate interface. A DHCP server dynamically assigns IP addresses to hosts on the network connected to the interface. The host computers must be configured to obtain their IP addresses using DHCP.

If an interface is connected to multiple networks through routers, you can add a DHCP server for each network. The IP range of each DHCP server must match the network address range. The routers must be configured for DHCP relay.

You can configure a FortiGate interface as a DHCP relay. The interface forwards DHCP requests from DHCP clients to an external DHCP server and returns the responses to the DHCP clients. The DHCP server must have appropriate routing so that its response packets to the DHCP clients arrive at the unit.

DHCP server options aren't available in transparent mode.

Configuring DHCP servers

To add a DHCP server, go to **Network > Interfaces**. Edit the interface, and select **DHCP** in the addressing mode.

Field	Description
Address Range	By default, the FortiGate unit assigns an address range based on the address of the interface for the complete scope of the address. For example, if the interface address is 172.20.120.230, the default range created is 172.20.120.231 to 172.20.120.254. Select the range and select Edit to adjust the range or select Create New to add a different range.
Netmask	Enter the netmask of the addresses that the DHCP server assigns.
Default Gateway	Select this to use either Same as Interface IP or select Specify and enter the IP address of the default gateway that the DHCP server assigns to DHCP clients.
DNS Server	Select this to use Same as system DNS , Same as Interface IP or select Specify and enter the IP address of the DNS server.
Mode	Select the type of DHCP server FortiGate will be. By default, it is a Server . Select Relay if needed. When Relay is selected, the above configuration is replaced by a field to enter the DHCP Server IP address.
DHCP Server IP	This appears only when Mode is Relay . Enter the IP address of the DHCP server where FortiGate obtains the requested IP address.

Field	Description
Type	Select this to use the DHCP in Regular or IPsec mode.
Additional DHCP Options	Use this to create new DHCP options.
MAC Address + Access Control	Select this to match an IP address from the DHCP server to a specific client or device using its MAC address. In a typical situation, an IP address is assigned ad hoc to a client, and that assignment times out after a specific time of inactivity from the client, known as the lease time. To ensure a client or device always has the same IP address (there is no lease time), use IP reservation.
Add from DHCP Client List	If the client is currently connected and using an IP address from the DHCP server, you can select this option to select the client from the list.

DHCP Server

 Advanced...

Mode

Server

Relay

DHCP Server IP

Type

Regular

IPsec


```
edit <name>
    set dhcp-relay-agent-option {enable | disable}
next
end
```

For more information about the DHCP relay option, see [RFC 3046](#) (DHCP Relay Agent Information Option).

Configuring DHCP with IPv6

You can use DHCP with IPv6, using the CLI. To configure DHCP, ensure IPv6 is enabled by going to **System > Feature Visibility** and enable **IPv6** under **Basic Features**. Use the following CLI command:

```
config system dhcp6 server
```

For more information about the configuration options, see the [FortiOS CLI Reference](#).

DHCPv6 prefix delegation

FortiGate supports prefix delegation for DHCP for IPv6 addressing. It's not practical to manually provision networks on a large scale in IPv6 networking. You can use DHCPv6 prefix delegation to assign a network address prefix, and automate the configuration and provisioning of the public routable addresses for the network.

To enable the prefix delegation - CLI:

```
config system interface
    edit "wan1"
        config ipv6
            set ip6-mode dhcp
            set ip6-allowaccess ping
            set dhcp6-prefix-delegation enable
        next
    next
end
```

Range for DHCPv6 prefix delegation

You can configure a range for DHCPv6 server prefix delegation. You can add a prefix range (starting and ending prefixes) and a prefix length. The prefix length determines the length of the prefix that the FortiGate sends downstream.

To configure a range for DHCPv6 prefix delegation – CLI:

```
config system dhcp6 server
    edit <id>
        config prefix-range
            edit <id>
                set start-prefix <prefix>
                set end-prefix <prefix>
                set prefix-length <length>
            next
        next
    next
end
```

DHCPv6 prefix hint

This feature is used to "hint" to upstream DHCPv6 servers a desired prefix length for their subnet to be assigned in response to its request.

There is a possibility of duplicate prefixes being sent by ISP when using a /64 bit subnet because the first 64 bits of the address are derived from the MAC address of the interface. This could cause an issue if the system administrator wishes to divide the host networks into 2 /64 bit subnets.

By receiving a /60 bit (for example) network address, the administrator can then divide the internal host works without the danger of creating duplicate subnets.

Also included in the new feature, are preferred times for the life and valid life of the DHCP lease.

DHCPv6 hint for the prefix length:

```
set dhcp6-prefix-hint <DHCPv6 prefix that will be used as a hint to the upstream DHCPv6 server>
```

DHCPv6 hint for the preferred life time:

```
set dhcp6-prefix-hint-plt <integer> 1 ~ 4294967295 seconds or "0" for unlimited lease time
```

DHCPv6 hint for the valid life time:

```
set dhcp6-prefix-hint-vlt <integer> 1 ~ 4294967295 seconds or "0" for unlimited lease time
```

Service

On low-end FortiGate units, a DHCP server is configured on the internal interface, by default, with the following values:

Field	Value
Address Range	192.168.1.110 to 192.168.1.210
Netmask	255.255.255.0
Default Gateway	192.168.1.99
Lease Time	7 days
DNS Server 1	192.168.1.99

These settings are appropriate for the default internal interface IP address of 192.168.1.99. If you change this address to a different network, you need to change the DHCP server settings to match.

Alternatively, after the FortiGate unit assigns an address, you can go to **Monitor > DHCP Monitor** and locate the specific user. Right-click and select **Create/Edit IP Reservation**.

Configuring the lease time

The lease time determines the length of time an IP address remains assigned to a client. Once the lease expires, the address is released for allocation to the next client that requests an IP address.

To configure the lease time, use the following CLI commands:

```
config system dhcp server
```

```
edit <server_entry_number>
    set lease-time <seconds>
next
end
```

The default lease time is seven days. To have an unlimited lease time, set the value to zero.

Configuring TFTP servers

You can configure multiple Trivial File Transfer Protocol (TFTP) servers for a Dynamic Host Configuration Protocol (DHCP) server. For example, you may want to configure a main TFTP server and a backup TFTP server.

The `tftp-server` command allows you to configure the TFTP servers, using either their hostnames or IP addresses. Separate multiple server entries with spaces.

To configure TFTP servers - CLI:

```
config system dhcp server
    edit <server ID>
        set tftp-server <hostname/IP address> <hostname/IP address>
    next
end
```

Configuring the DHCP renew time

You can set a minimum DHCP renew time. This option is available only when `mode` is set to `dhcp`.

To set the DHCP renew time - CLI:

```
config system interface
    edit <name>
        set mode dhcp
        set dhcp-renew-time <seconds>
    next
end
```

The possible values for `dhcp-renew-time` are 300 to 605800 seconds (five minutes to seven days). To use the renew time that the server provides, set this entry to 0.

DHCP options

When you add a DHCP server, you can include DHCP codes and options. The DHCP options are BOOTP vendor information fields that provide additional vendor-independent configuration parameters to manage the DHCP server. For example, you may need to configure a FortiGate DHCP server that gives out a separate option, as well as an IP address, such as an environment that needs to support PXE boot with Windows images.

The option numbers and codes are specific to a particular application. The documentation for the application should provide the values you should use. Option codes are represented in option value and HEX value pairs. The option is a value between 1 and 255.

You can add up to three DHCP code/option pairs per DHCP server.

To configure option 252 with value `http://192.168.1.1/wpad.dat` - CLI:

```
config system dhcp server
    edit <server_entry_number>
```

```

        set option1 252 687474703a2f2f3139322e3136382e312e312f777061642e646174
    next
end

```

For more information about DHCP options, see [RFC 2132](#) (DHCP Options and BOOTP Vendor Extensions).

FortiGate DHCP works with DDNS to allow FQDN connectivity to leased IP addresses

As clients are assigned IP addresses, they send back information that would be found in an A record to the FortiGate DHCP server, which can take this information and pass it back to a corporate DNS server so that even devices using leased IP address can be reached using FQDNs. You can configure the settings for this feature using the `ddns-update` CLI command and some other `ddns` related options.

DHCP server option fields

In place of specific fields, the DHCP server maintains a table for the potential options. The FortiOS DHCP server supports up to a maximum of 30 custom options. These optional fields are set in the CLI.

To get to the DHCP server - CLI:

```

config system dhcp server
    edit <integer - ID of the specific DHCP server>

```

To configure the options, use the following CLI command:

```

config options

```

Once you are in the options context, create an ID for the table entry, using the following CLI commands:

```

edit <integer>
    set code <integer between 0 - 4294967295 to determine the DHCP option>
    set type [ hex | string | ip ]
    set value <option content for DHCP option types hex and string>
    set ip <option content for DHCP option type ip>
end

```

Excluding addresses in DHCP

If you have a large address range for the DHCP server, you can block a range of addresses that won't be included in the available addresses for the connecting users.

To exclude addresses in DHCP - CLI:

```

config system dhcp server
    edit <server_entry_number>
        config exclude-range
            edit <sequence_number>
                set start-ip <address>
                set end-ip <address>
            next
        next
    next
end

```


Viewing information about DHCP server connections

To view information about DHCP server connections, go to **Monitor > DHCP Monitor**. On this page, you can also add IP addresses to the reserved IP address list.

Breaking an address lease

If you need to end an IP address lease, you can break the lease. This is useful if you have limited addresses and longer lease times when some leases are no longer necessary, for example, with corporate visitors.

To break a lease - CLI:

```
execute dhcp lease-clear <ip_address>
```

Interface MTU packet size

You can change the maximum transmission unit (MTU) of the packets that FortiGate transmits to improve network performance. Ideally, the MTU should be the same as the smallest MTU of all the networks between FortiGate and the destination of the packets. If the packets that the FortiGate sends are larger than the smallest MTU, they're broken up or fragmented, which slows down transmission. You can easily experiment by lowering the MTU to find an MTU size for optimum network performance.

- 68 to 1500 bytes for static mode
- 576 to 1500 bytes for DHCP mode
- 576 to 1492 bytes for PPPoE mode
- Larger frame sizes (if supported by the FortiGate model), up to 9216 bytes for NP2, NP4, and NP6-accelerated interfaces

This option is available only for physical interfaces. Virtual interfaces associated with a physical interface inherit the physical interface MTU size.

Interfaces on some FortiGate models support frames larger than the traditional 1500 bytes. Jumbo frames are supported on FortiGate models that have either a SOC2 or NP4lite (except for the FortiGate 30D), and on FortiGate 100D series models. For information about your FortiGate model's hardware, see the [FortiOS Hardware Acceleration Handbook](#). To find out the maximum frame size that's supported for other models, visit the [Fortinet Support](#) website.

If you need to send larger frames over a route, all Ethernet devices on that route must support the larger frame size. Otherwise, the larger frames won't be recognized and will be dropped.

If you have standard size and larger size frame traffic on the same interface, routing alone can't route them to different routes based only on frame size. However, you can use VLANs to make sure the larger frame traffic is routed over network devices that support the larger size. VLANs inherit the MTU size from the parent interface. You must configure the VLAN to include both ends of the route, as well as all switches and routers along the route.

You can configure the MTU packet size. If you select an MTU size larger than your FortiGate model supports, an error message will indicate this. In this situation, try configuring a smaller MTU size until the value is supported.



In transparent mode, if you change the MTU of an interface, you must change the MTU of all interfaces on FortiGate to match the new MTU.

To change the MTU size - CLI:

```

config system interface
  edit <interface name>
    set mtu-override enable
    set mtu <byte size>
  next
end

```

Interface settings

You configure FortiGate interfaces, both physical and virtual, in **Network > Interfaces** in the FortiGate GUI. There are different options for configuring interfaces when FortiGate is in NAT mode or transparent mode.

On FortiOS Carrier, you can also enable the Gi gatekeeper on each interface for anti-overbilling.

Field	Description
Create New	<p>Select this to add a new interface, zone, or virtual wire pair.</p> <p>Depending on the FortiGate model, you can add a VLAN interface, a loopback interface, an IEEE 802.3ad aggregated interface, or a redundant interface.</p> <p>When VDOMs are enabled, you can also add Inter-VDOM links.</p>
Name	<p>The names of the physical interfaces on FortiGate. This includes any alias names that have been configured.</p> <p>When you combine several interfaces into an aggregate or redundant interface, only the aggregate or redundant interface is listed, and not the component interfaces.</p> <p>If you added VLAN interfaces, they appear in the name list below the physical or aggregated interface to which they have been added.</p> <p>If you added loopback interfaces, they appear in the interface list below the physical interface to which they have been added. If software switch interfaces are configured, you can view them.</p> <p>If your FortiGate model supports AMC modules, the interfaces are named amc-sw1/1, amc-dw1/2, and so on.</p>
Type	The configuration type for the interface.
IP/Netmask	<p>The current IP address and netmask of the interface.</p> <p>In VDOM, when VDOMs are not all in NAT or transparent mode, some values may not be available for display and are displayed as “-”.</p>
Access	The administrative access configuration for the interface.

Field	Description
Administrative Status	<p>Indicates if the interface can be accessed for administrative purposes. If the administrative status is a green arrow, an administrator can connect to the interface using the configured access.</p> <p>If the administrative status is a red arrow, the interface is administratively down and can't be accessed for administrative purposes.</p>
Link Status	<p>The status of the interface physical connection. The link status can be up (green arrow) or down (red arrow). If the link status is up, the interface is connected to the network and accepting traffic. If the link status is down, the interface is either not connected to the network or there is a problem with the connection.</p> <p>You can't change the link status from the FortiGate GUI, and it typically indicates that an Ethernet cable is plugged into the interface.</p> <p>The link status is only displayed for physical interfaces.</p>
MAC	The MAC address of the interface.
Mode	The addressing mode of the interface. This value can be manual, DHCP, or PPPoE.
Secondary IP	The secondary IP addresses added to the interface.
MTU	The maximum number of bytes per transmission unit for the interface.
Virtual Domain	The virtual domain to which the interface belongs. This column is visible when VDOM configuration is enabled.
VLAN ID	The configured VLAN ID for VLAN subinterfaces.

Interface configuration and settings

To configure an interface, go to **Network > Interfaces**, and select **Create New**.

Name	Enter the name of the interface. Physical interface names can't be changed.
Alias	<p>Enter an alternate name for a physical interface on the FortiGate unit. This field appears when you edit an existing physical interface.</p> <p>The alias is a maximum of 25 characters. The alias name doesn't appear in logs.</p>
Link Status	Indicates whether the interface is connected to a network (link status is Up) or not (link status is Down). This field appears when you edit an existing physical interface.

Type	<p>Select the type of interface you want to add.</p> <p>On some FortiGate models, you can set Type to 802.3ad Aggregate or Redundant Interface.</p>
Interface	<p>This is displayed when Type is set to VLAN.</p> <p>Select the name of the physical interface that you want to add a VLAN interface to. Once created, the VLAN interface is listed below its physical interface in the Interface list.</p> <p>You can't change the physical interface of a VLAN interface except when you add a new VLAN interface.</p>
VLAN ID	<p>This is displayed when Type is set to VLAN.</p> <p>Enter the VLAN ID. You can't change the VLAN ID except when you add a new VLAN interface.</p> <p>The VLAN ID must be a number between 1 and 4094. It must match the VLAN ID that the IEEE 802.1Q-compliant router or switch that is connected to the VLAN subinterface adds.</p>
Virtual Domain	<p>Select the virtual domain to add the interface to.</p> <p>Administrator accounts with the super_admin profile can change the Virtual Domain.</p>
Physical Interface Members	<p>This section can have two different formats depending on the interface type:</p> <ul style="list-style-type: none"> • Software switch interface: This section is a display-only field that shows the interfaces that belong to the virtual interface of the software switch. • 802.3ad aggregate or Redundant interface: This section includes the available interface list and the selected interface list. <p>Select interfaces from the Available Interfaces list and select the right arrow to add an interface to the Selected Interface list.</p>

Addressing mode	<p>Select the addressing mode for the interface:</p> <ul style="list-style-type: none"> • Select Manual and add an IP/Netmask for the interface. If IPv6 configuration is enabled, you can add both a IPv4 and an IPv6 IP address. • Select DHCP to get the interface IP address and other network settings from a DHCP server. • Select PPPoE to get the interface IP address and other network settings from a PPPoE server. • Select One-Arm Sniffer to enable the interface as a means to detect possible traffic threats. This option is available on physical ports that aren't configured for the primary Internet connection. • Select Dedicate to FortiAP/FortiSwitch to have a FortiAP or FortiSwitch device connect exclusively to the interface. This option is available only when you edit a physical interface and it has a static IP address. When you enter the IP address, FortiGate automatically creates a DHCP server using the subnet that you enter. This option is not available on the ADSL interface. <p>The FortiSwitch option is currently available only on the FortiGate 100D.</p>
IP/Netmask	If Addressing Mode is set to Manual , enter an IPv4 address and subnet mask for the interface. FortiGate interfaces can't have IP addresses on the same subnet.
IPv6 Address	If Addressing Mode is set to Manual and IPv6 support is enabled, enter an IPv6 address and subnet mask for the interface. A single interface can have an IPv4 address, IPv6 address, or both.
Administrative Access	Select the types of administrative access that you want to allow for IPv4 connections to this interface.
HTTPS	Allow secure HTTPS connections to the FortiGate GUI through this interface. If configured, this option will enable automatically when you select the HTTP option.
PING	The interface responds to pings. Use this setting to verify your installation and for testing.
HTTP	Allow HTTP connections to the FortiGate GUI through this interface. If configured, this option will also enable the HTTPS option.
SSH	Allow SSH connections to the CLI through this interface.
SNMP	Allow a remote SNMP manager to request SNMP information by connecting to this interface.
FMG-Access	Allow FortiManager authorization automatically during the communication exchanges between FortiManager and FortiGate devices.

CAPWAP	Allows the FortiGate wireless controller to manage a wireless access point, such as a FortiAP device.
IPv6 Administrative Access	Select the types of administrative access that you want to allow for IPv6 connections to this interface. The types are the same as for Administrative Access.
Security Mode	Select a captive portal for the interface. After you select this, you can define the portal message and the appearance of the GUI that users see when they log into the interface. You can also define one or more user groups that can access the interface.
DHCP Server	Select this to enable a DHCP server for the interface. For more information about configuring a DHCP server on the interface, see "DHCP servers and relays" on page 2002 .
Device Detection	Select this to allow the interface to be used with BYOD devices, such as iPhones. Define the device definitions by selecting User & Device > Device Inventory in the FortiGate GUI.
Enable Explicit Web Proxy	<p>Select this to enable explicit web proxying on this interface.</p> <p>This is available when you enable explicit proxy in the System Information Dashboard (System > Dashboard > Status).</p> <p>When you enable this, the interface will be displayed in System > Network > Explicit Proxy, under Listen on Interfaces, and web traffic on this interface will be proxied according to the Web Proxy settings.</p> <p>This option isn't available for a VLAN interface selection.</p>
Secondary IP Address	Add additional IPv4 addresses to this interface. Select the expand arrow to expand or hide the section.
Comments	Enter a description (up to 63 characters) to describe the interface.
Gi Gatekeeper (FortiOS Carrier only)	For FortiOS Carrier, enable this to enable the Gi firewall as part of the anti-overbilling configuration. You must also configure Gi Gatekeeper Settings by selecting System > Admin > Settings in the FortiGate GUI.

Interface Name wan2 (08:5B:0E:50:9D:B2)

Alias

Link Status Down

Type Physical Interface

Role Undefined

Address

Addressing mode Manual **DHCP** PPPoE One-Arm Sniffer Dedicated to FortiSwitch

Status initializing.....

Retrieve default gateway from server ☐

Override internal DNS ☒

IPv6 Addressing mode **Manual** DHCP

IPv6 Address/Prefix

Restrict Access

Administrative Access	<input type="checkbox"/> HTTPS	<input checked="" type="checkbox"/> PING	<input checked="" type="checkbox"/> FMG-Access	<input type="checkbox"/> CAPWAP	<input type="checkbox"/> SSH
	<input type="checkbox"/> SNMP	<input type="checkbox"/> RADIUS Accounting	<input type="checkbox"/> FortiTelemetry		
IPv6 Administrative Access	<input type="checkbox"/> HTTPS	<input type="checkbox"/> PING	<input type="checkbox"/> FMG-Access	<input type="checkbox"/> CAPWAP	<input type="checkbox"/> SSH
	<input type="checkbox"/> SNMP				

Networked Devices

Device Detection ☐

Miscellaneous

Scan Outgoing Connections to Botnet Sites **Disable** Block Monitor

Enable Explicit Web Proxy ☐

Status

Comments 0/255

Interface State **Enabled** Disabled

OK Cancel

If you assign an interface to be part of a virtual wire pairing, the "role" value is removed from the interface.

Loopback interfaces

A loopback interface is a logical interface that's always up (no physical link dependency) and the attached subnet is always present in the routing table.

The IP address of the FortiGate loopback interface doesn't depend on one specific external port, and therefore you can access it through several physical or VLAN interfaces. You can configure multiple loopback interfaces in either non-VDOM mode or in each VDOM.

Loopback interfaces still require appropriate firewall policies to allow traffic to and from the interfaces.

A loopback interface can be used with:

- Management access
- BGP (TCP) peering
- PIM RP
- IS-IS

Loopback interfaces are a good practice for OSPF. To make troubleshooting OSPF easier, you should set the OSPF router ID to the same value as the loopback IP address, and remember the management IP addresses (ssh to “router ID”).

You can enable dynamic routing protocols on loopback interfaces.

For blackhole static routes, use the blackhole route type instead of the loopback interface.

VXLAN loopback binding

A Virtual Extensible LAN (VXLAN) unicast device can bind to a loopback interface as its underlying interface. The IP address of the loopback interface is taken as the source IP address for its outgoing VXLAN packets so the peer knows where to reply. Among the parameters that are passed to the kernel, the ifindex of the loopback interface isn't passed down to the kernel, so the kernel can choose the outgoing physical interface. This way, VXLAN traffic can be routed across multiple physical links and it provides resistance to a single point of failure.

To configure VXLAN loopback binding - CLI:

```
config system vxlan
  edit <name>
    set interface <interface>
    set vni <VXLAN network ID>
    set remote-ip <IP address>
  next
end
```

One-armed sniffer

You can use a one-armed sniffer to configure a FortiGate physical interface as a one-arm intrusion detection system (IDS). Traffic sent to the interface is examined for matches to the configured IPS sensor and application control list. Matches are logged and then all received traffic is dropped. Sniffing reports only on attacks. It does not deny or otherwise influence traffic.

You can use the one-arm sniffer to configure FortiGate to operate as an IDS appliance by sniffing network traffic for attacks without actually processing the packets. To configure one-arm IDS, you enable sniffer mode on a FortiGate interface and connect the interface to a hub, or to the SPAN port of a switch that is processing network traffic.


To assign an interface as a sniffer interface, select **Network > Interfaces**, edit the interface and select **One-Arm Sniffer**.

If the check box is not available, it means the interface is in use. Ensure that the interface isn't selected in any firewall policies, routes, virtual IPs, or other features in which a physical interface is specified.



Field	Description
Enable Filters	<p>Select this to include filters that define a more granular sniff of network traffic. Select specific hosts, ports, VLANs, and protocols.</p> <p>In all cases, enter a number or number range for the filtering type. For protocol values, the standard protocols are:</p> <ul style="list-style-type: none">• UDP - 17• TCP - 6• ICMP - 1
Include IPv6 Packets	If your network is running both IPv4 and IPv6 addressing, select this to sniff both addressing types. Otherwise, FortiGate will sniff only IPv4 traffic.
Include Non-IP Packets	Select this for a more intense scan of content in the traffic.
Security Profiles	IPS sensors and application control lists allow you to select specific sensors and applications that you want to identify within the traffic.

Interface Name wan2 (08:5B:0E:50:9D:B2)

Alias

Link Status Down 


Type Physical Interface


Role  Undefined 


Address


Addressing mode Manual DHCP PPPoE One-Arm Sniffer Dedicated to FortiSwitch

☒ Enable Filters

Host(s) 

Port(s) 


VLAN(s) 


Protocol 


☐ Include IPv6 Packets


☐ Include Non-IP Packets


Security Profiles

Enable AntiVirus  [Edit Sniffer Profile](#)


Enable Web Filter  [Edit Sniffer Profile](#)

Enable Application Control  [Edit Sniffer Profile](#)

Enable CASI Profile  [Edit Sniffer Profile](#)

Enable IPS  [Edit Sniffer Profile](#)

Logging Options

Log Allowed Traffic  Security Events All Sessions

Scan Outgoing Connections to Botnet Sites Disable Block Monitor

Ports preassigned as sniffer ports

Some FortiGate models have ports preconfigured as sniffer ports, by default. The models and ports preconfigured in sniffer mode are as follows:

- FortiGate 300D
 - Port4
 - Port8
- FortiGate 500D
 - Port5
 - Port6
 - Port13
 - Port14


Physical ports

FortiGate has several physical ports that you can connect Ethernet or optical cables to. Depending on the FortiGate model, it can have between 4 and 40 physical ports. Some units have a grouping of ports labeled as **lan**, that provide built-in switch functionality.

The port names, as labeled on the FortiGate, appear in the FortiGate GUI in the **Unit Operation** widget on the dashboard. They also appear when you configure the interfaces, in **Network > Interfaces**.

You can hover over the ports to see information about each port, such as the name of the port and the IP address.

For example, the following diagram shows the 22 interfaces of the FortiGate 100 D (Generation 2) as they appear in the dashboard in the FortiGate GUI.



Two of the physical ports on the FortiGate 100D (Generation 2) are SFP ports. These ports share the numbers 15 and 16 with RJ-45 ports. Because of this, when SFP port 15 is used, RJ-45 port 15 can't be used, and vice versa. These ports also share the same MAC address.

Configuring the FortiGate 100D ports

Normally, you can configure the internal interface as a single interface that's shared by all physical interface connections (a switch). The switch mode feature has two states: switch mode and interface mode. Switch mode is the default mode, with only one interface and one address for the entire internal switch. Interface mode allows you to configure each of the physical interface connections of the internal switch separately. This allows you to assign different subnets and netmasks to each of the internal physical interface connections.

The larger FortiGate models may also include Advanced Mezzanine Cards (AMC), which can provide additional interfaces (Ethernet or optical), with throughput enhancements for more efficient handling of specialized traffic. These interfaces appear in FortiOS as port `amc/sw1`, `amc/sw2`, and so on.

Displaying information about the status of transceivers

You can display information about the status of transceivers installed in FortiGate SFP/SFP+ interfaces, in the FortiGate CLI.

The `get system interface transceiver` command lists all of the SFP/SFP+ interfaces on FortiGate. If the interfaces include transceivers, the command output displays information about them, such as the vendor name, part number, and serial number. It also includes details about transceiver operation, such as temperature, voltage, and optical transmission power, which you can use to diagnose transmission problems.

The following example shows an output from using this command:

```
get system interface transceiver
...
Interface port14 - Transceiver is not detected.
Interface port15 - SFP/SFP+
```

```

Vendor Name   : FIBERXON INC.
Part No.      : FTM-8012C-SLG
Serial No.    : 101680071708917
Interface port16 - SFP/SFP+
Vendor Name   : FINISAR CORP.
Part No.      : FCLF-8521-3
Serial No.    : PS62ENQ

```

SFP/SFP+ Interface	Temperature (Celsius)	Voltage (Volts)	Optical Tx Bias (mA)	Optical Tx Power (dBm)	Optical Rx Power (dBm)
port15	N/A	N/A	N/A	N/A	N/A
port16	N/A	N/A	N/A	N/A	N/A

++ : high alarm, + : high warning, - : low warning, -- : low alarm, ? : suspect.

You can use this command on most FortiGate models that have SFP/SFP+ interfaces.

Split port support

The 5001D 40 GB can be split into 4 10 GB ports. You can do this through a combination of hardware and software configuration. You use a specific 40 GB connector to connect to the 40 GB port and typically, the other end of the fibre optic cable connects to another 40 GB port. However, you can use a special cable that is a single 40 GB connector at one end and 4 10 GB connections at the other end. To use this setup, you also have to configure the port to be a split port.

To configure split port support - CLI:

```

config system global
    set port-split port1 port2
end

```

The ports will be checked to make sure that they aren't in use or referenced by other policy configurations. If they are in use, the command is aborted. Changing the port to be a split port requires a system reboot.

PPPoE addressing mode on an interface

If you configure the interface to use PPPoE, the FortiGate automatically broadcasts a PPPoE request from the interface.

FortiGate devices support many PPPoE RFC features (RFC 2516) including unnumbered IPs, initial discovery timeout and PPPoE Active Discovery Terminate (PADT).

PPPoE is only configurable in the GUI on desktop FortiGate devices. 1U FortiGate devices and up must be configured in the CLI.

To configure PPPoE - CLI:

```

config system interface
    edit <port name>
        set mode pppoe
        set username <ISP username>
        set password <ISP password>
        set idle-timeout <seconds>
        set distance <integer>
    end
end

```

```

    set ipunnumbered <unnumbered IP>
    set disc-retry-timeout <seconds>
    set padt-retry-timeout <seconds>
    set lcp-echo-interval <seconds>
    set dns-server-override {enable | disable}
  next
end

```

To configure PPPoE - GUI:

Configure PPPoE on an interface in **Network > Interfaces**. The following table describes the PPPoE status information when PPPoE is configured for an interface.

Field	Description
Status	<p>Displays PPPoE status messages as the FortiGate connects to the PPPoE server and gets addressing information. Select Status to refresh the addressing mode status message.</p> <p>The status is only displayed if you selected Edit.</p> <p>Status can be any one of the following 4 messages.</p>
Initializing	No activity.
Connecting	The interface is attempting to connect to the PPPoE server.
Connected	<p>The interface retrieves an IP address, netmask, and other settings from the PPPoE server.</p> <p>When the status is connected, PPPoE connection information is displayed.</p>
Failed	The interface was unable to retrieve an IP address and other information from the PPPoE server.
Reconnect	<p>Select to reconnect to the PPPoE server.</p> <p>Only displayed if Status is connected.</p>
User Name	The username for the PPPoE account.
Password	The password for the PPPoE account.
Unnumbered IP	Specify the IP address for the interface. If your ISP has assigned you a block of IP addresses, use one of them. Otherwise, this IP address can be the same as the IP address of another interface or can be any IP address.
Initial Disc Timeout	Enter Initial discovery timeout. Enter the time to wait before starting to retry a PPPoE discovery.

Field	Description
Initial PADT timeout	Enter Initial PPPoE Active Discovery Terminate (PADT) timeout, in seconds. Use this timeout to shut down the PPPoE session if it's idle for the specified number of seconds. PADT must be supported by your ISP. Set the Initial PADT timeout to 0 to disable.
Distance	Enter the administrative distance for the default gateway retrieved from the PPPoE server. The administrative distance, an integer from 1-255, specifies the relative priority of a route when there are multiple routes to the same destination. A lower administrative distance indicates a more preferred route. The default distance for the default gateway is 1.
Retrieve default gateway from server	Enable to retrieve a default gateway IP address from a PPPoE server. The default gateway is added to the static routing table.
Override internal DNS	<p>Enable to replace the DNS server IP addresses on the System DNS page with the DNS addresses retrieved from the PPPoE server.</p> <p>When VDOMs are enabled, you can override the internal DNS only on the management VDOM.</p>

Probing interfaces

You can use server probes on FortiGate interfaces. First, you configure the probe response mode and then you give the probe response administrative access on the interface. You can configure one of the following probe response modes:

Mode	Description
none	Disable probe
http-probe	HTTP probe
twamp	Two-Way Active Measurement Protocol

To configure the probe - CLI:

```
config system probe-response
    set mode {none|http-probe|twamp}
end
```

To give the probe response administrative access to the interface - CLI:

```
config system interface
    edit <port>
        set allowaccess probe-response
    next
end
```

TWAMP

FortiOS supports Two-Way Active Measurement Protocol (TWAMP) Light, which is a simplified architecture within the TWAMP standard. Its purpose is to measure the round trip IP performance between any two devices within a network that supports the protocol. FortiOS supports both responder/reflector and server/controller roles.

FortiOS extends TWAMP to also support unidirectional network quality monitoring. You can monitor network quality for each direction of a traffic path separately. You can also use SNMP to monitor the network quality status from both the controller and responder sides.

You can use a link health monitor to see the following information:

- Directional latency (Minimum, Maximum, Average)
- Directional packet loss
- Directional out of sequence packets
- Directional jitter (Minimum, Maximum, Average)

To use TWAMP to monitor network quality - CLI:

```
config system link-monitor
  edit <name>
    set srcintf <interface name>
    set server <IP address>
    set protocol twamp
    set port <port number>
    set gateway-ip <IP address>
  next
end
```

To configure the probe to use TWAMP – CLI:

```
config system probe-response
  set mode twamp
  set timeout <time>
end
```

Redundant interfaces

On some models, you can combine two or more physical interfaces to provide link redundancy. This feature allows you to connect to two or more switches to ensure connectivity if one physical interface, or the equipment on that interface, fails.

In a redundant interface, traffic travels only over one interface at a time. This differs from an aggregated interface where traffic travels over all interfaces for distribution of increased bandwidth. This difference means that redundant interfaces can have more robust configurations with fewer possible points of failure. This is important in a fully-meshed HA configuration.

An interface can be in a redundant interface if:

- It's a physical interface, not a VLAN interface
- It's not already part of an aggregated or redundant interface
- It's in the same VDOM as the redundant interface
- It has no defined IP address
- It's not configured for DHCP or PPPoE

- It has no DHCP server or relay configured on it
- It doesn't have any VLAN subinterfaces
- It isn't referenced in any security policy, VIP, or multicast policy
- It isn't monitored by HA
- It isn't one of the FortiGate-5000 series backplane interfaces

When an interface is included in a redundant interface, it isn't listed on the **Network > Interfaces** page. You can't configure the interface individually and it isn't available for inclusion in security policies, VIPs, or routing.

Dual Internet connections

Dual internet connections, also referred to as dual WAN or redundant Internet connections, refers to using two FortiGate interfaces to connect to the Internet. You can use dual Internet connections in several ways:

- Redundant interfaces: If one interface goes down, the second interface automatically becomes the main Internet connection.
- Load sharing: This ensures better throughput.
- Use a combination of redundancy and load sharing.

Redundant interfaces

Redundant interfaces ensure that if your Internet access is no longer available through a certain port, the FortiGate uses an alternate port to connect to the Internet.

In this scenario, two interfaces, WAN1 and WAN2, are connected to the Internet using two different ISPs. WAN1 is the primary connection. In the event of a failure of WAN1, WAN2 automatically becomes the connection to the Internet. For this configuration to function correctly, you must configure the following settings:

- Configure a link health monitor to determine when the primary interface (WAN1) is down and when the connection returns.
- Configure a default route for each interface.
- Configure security policies to allow traffic through each interface to the internal network.

Link health monitor

Adding a link health monitor is required for routing failover traffic. A link health monitor confirms the connectivity of the device's interface. You can detect possible routing loops with link health monitors. You can configure the FortiGate to ping a gateway at regular intervals to ensure it's online and working. When the gateway isn't accessible, that interface is marked as down. After this configuration, when this interface on the FortiGate can't connect to the next router, the FortiGate brings down the interface.

Set the `Interval` (how often to send a ping) and `failtime` (how many lost pings are considered a failure). A smaller interval and smaller number of lost pings results in faster detection, but creates more traffic on your network. You may also want to log CPU and memory usage, as a network outage will cause your CPU activity to spike.

The link health monitor supports both IPv4 and IPv6. For IPv6, it supports the ping6 protocol only.

To add a link health monitor (IPv4) - CLI:

```
config system link-monitor
  edit <link-monitor-name>
    set addr-mode ipv4
```



```
set srcint <interface-name>
set server <server-IP-address>
set protocol {ping tcp-echo udp-echo http twamp}
set gateway-ip <gateway-IP-address>
set source-ip <IP-address>
set interval <seconds>
set failtime <retry-attempts>
set recoverytime <number-of-successful-responses>
set ha-priority <priority>
set update-cascade-interface {enable | disable}
set update-static-route {enable | disable}
set status enable
next
end
```

To add a link health monitor (IPv6) - CLI:

```
config system link-monitor
edit <link-monitor-name>
set addr-mode ipv6
set srcint <interface-name>
set server <server-IP-address>
set protocol ping6
set gateway-ip6 <gateway-IP-address>
set source-ip6 <IP-address>
set interval <seconds>
set failtime <retry-attempts>
set recoverytime <number-of-successful-responses>
set ha-priority <priority>
set update-cascade-interface {enable | disable}
set update-static-route {enable | disable}
set status enable
next
end
```

Routing

You must configure a default route for each interface and indicate which route is preferred by specifying the distance. The lower distance is declared active and placed higher in the routing table.



When you have dual WAN interfaces that are configured to provide failover, you might not be able to connect to the backup WAN interface because the FortiGate may not route traffic (even responses) out of the backup interface. The FortiGate performs a reverse path lookup to prevent spoofed traffic. If an entry can't be found in the routing table that sends the return traffic out the same interface, the incoming traffic is dropped.

To configure the routing of the two interfaces - GUI:

1. Go to **Network > Static Routes** and select **Create New**.
2. Enter the following information and select **OK**.

Destination IP/Mask	For an IPv4 route, enter a subnet of 0.0.0.0/0.0.0.0. For an IPv6 route, enter a subnet of ::/0.
Gateway	Enter the gateway address.
Interface	Select the primary connection. For example, WAN1 .
Administrative Distance	Leave as the default of 10.

- Repeat these steps to set **Interface** to **WAN2** and **Administrative Distance** to 20.

To configure the routing of the two interfaces - CLI:

```
config router {static | static6}
  edit 0
    set dst 0.0.0.0 0.0.0.0
    set device WAN1
    set gateway <gateway_address>
    set distance 10
  next
  edit 0
    set dst 0.0.0.0 0.0.0.0
    set device WAN2
    set gateway <gateway_address>
    set distance 20
  next
end
```

Security policies

When you create security policies, you need to configure duplicate policies to ensure that after traffic fails over WAN1, regular traffic is allowed to pass through WAN2, as it did with WAN1. This ensures that failover occurs with minimal effect to users. For more information about creating security policies, see the [FortiOS Firewall Handbook](#).

Load sharing

Load sharing allows you to use both connections to the Internet at the same time, but doesn't provide failover support. When configuring load sharing, you need to make sure that routing is configured for both external ports (for example, WAN1 and WAN2) have static routes with the same distance and priority.

Link redundancy and load sharing

In this scenario, both links are available to distribute Internet traffic over both links. Should one of the interfaces fail, the FortiGate will continue to send traffic over the other active interface. Configuration is similar to the Redundant interfaces configuration, with the main difference being that the configured routes should have equal distance settings.

This means both routes will remain active in the routing table. To make one interface the preferred interface, use a default policy route to indicate the interface that is preferred for accessing the Internet. If traffic matches the security policy, the policy overrides all entries in the routing table, including connected routes. You may need to add specific policy routes that override these default policy routes.

To redirect traffic over the secondary interface, create policy routes to direct some traffic onto it rather than the primary interface. When adding the policy route, only define the outgoing interface and leave the gateway blank. This ensures that the policy route won't be active when the link is down.

SSL VPN and WAN link load balancing

You can set virtual WAN link interfaces as destination interfaces in firewall policies for WAN link load balancing, when SSL VPN is the source interface. For example, you can log in to a FortiGate using an SSL VPN for traffic inspection and then have outbound traffic load balanced by WAN link load balancing.

You can set a virtual WAN link interface as a destination interface in a firewall policy where SSL VPN is the source interface, using either the FortiGate GUI (FortiOS 5.6.1 and later) or CLI.

To configure a virtual WAN link interface - CLI:

```
config firewall policy
  edit <policy ID>
    set dstintf virtual-wan-link
  next
end
```

Secondary IP addresses to an interface

If an interface is configured with a manual or static IP address, you can also add secondary static IP addresses to the interface. Adding secondary IP addresses effectively adds multiple IP addresses to the interface. Secondary IP addresses can't be assigned using DHCP or PPPoE.

All of the IP addresses added to an interface are associated with the single MAC address of the physical interface, and all secondary IP addresses are in the same VDOM as the interface that they're added to. You configure interface status detection for gateway load balancing separately for each secondary IP addresses. As with all other interface IP addresses, secondary IP addresses can't be on the same subnet as any other primary or secondary IP address assigned to a FortiGate interface unless they are in separate VDOMs.

To configure a secondary IP address, go to **Network > Interfaces**, select **Edit** or **Create New** and select the **Secondary IP Address** check box.

Software switch

A software switch, or soft switch, is a virtual switch that's implemented at the software, or firmware level, rather than the hardware level. A software switch can be used to simplify communication between devices connected to different FortiGate interfaces. For example, using a software switch, you can place the FortiGate interface connected to an internal network on the same subnet as your wireless interfaces. Then, devices on the internal network can communicate with devices on the wireless network without any additional configuration such as additional security policies, on the FortiGate.

It can also be useful if you require more hardware ports for the switch on a FortiGate. For example, if your FortiGate device has a 4-port switch, WAN1, WAN2 and DMZ interfaces, and you need one more port, you can create a soft switch that can include the 4-port switch and the DMZ interface all on the same subnet. These types of applications also apply to wireless interfaces and virtual wireless interfaces and physical interfaces, such as those with FortiWiFi and FortiAP devices.

Similar to a hardware switch, a software switch functions like a single interface. A software switch has one IP address; all of the interfaces in the software switch are on the same subnet. Traffic between devices connected to

each interface aren't regulated by security policies, and traffic passing in and out of the switch are affected by the same policy.

There are a few things to consider when setting up a software switch:

- Ensure you create a backup of the configuration.
- Ensure you have at least one port or connection such as the console port to connect to the FortiGate. If you accidentally combine too many ports, you will need a way to undo any errors.
- The ports that you include must not have any link or relation to any other aspect of the FortiGate. For example, DHCP servers, security policies, and so on.
- For increased security, you can create a captive portal for the switch, allowing only specific user groups access to the resources connected to the switch.
- To add an interface to a software switch, the interface can't be referenced by the existing configuration. It must also have its IP address set to 0.0.0.0/0.0.0.0.

To create a software switch - CLI:

```
config system switch-interface
    edit <switch-name>
        set type switch
        set member <interface_list>
    next
end
config system interface
    edit <switch_name>
        set ip <ip_address>
        set allowaccess https ssh ping
    next
end
```

Soft switch example

For this example, the wireless interface (Wi-Fi) needs to be on the same subnet as the DMZ1 interface to facilitate wireless syncing from an iPhone and a local computer. The syncing between two subnets is problematic. By putting both interfaces on the same subnet, the syncing will work. The software switch will accomplish this.



In this example, the soft switch includes a wireless interface. Remember to configure any wireless security before proceeding. If you leave this interface open without any password or other security, it leaves open access to not only the wireless interface but to any other interfaces and devices connected within the software switch.

Clear the interfaces and back up the configuration

First, ensure that the interfaces aren't being used with any other security policy or other use on the FortiGate. Check the Wi-Fi and DMZ1 ports to ensure that DHCP isn't enabled on the interface and there are no other dependencies with these interfaces.

Next, save the current configuration. In the event that something doesn't work, recovery can be quick.

Merge the interfaces

The plan is to merge the Wi-Fi port and DMZ1 port. This will create a software switch with a name of “synchro” with an IP address of 10.10.21.12. The following steps will create the switch, add the IP address and set administrative access for HTTPS, SSH, and Ping.

To merge the interfaces - CLI:

```
config system switch-interface
  edit synchro
    set type switch
    set member dmz1 wifi
  next
end
config system interface
  edit synchro
    set ip 10.10.21.12
    set allowaccess https ssh ping
  next
end
```

Final steps

With the switch set up, you can add security policies, DHCP servers, and any other configuration that you would normally do to configure interfaces on the FortiGate.

Virtual switch

The virtual switch feature allows you create virtual switches on top of physical switches with designated interfaces/ports so that a virtual switch can build up its forwarding table through learning and forward traffic accordingly. When traffic is forwarded among interfaces belonging to the same virtual switch, the traffic doesn't need to go up to the software stack, but is forwarded directly by the switch. When traffic has to be relayed to interfaces not on the virtual switch, the traffic will go through the normal data path and be offloaded to NP4, when possible.

This feature is only available on mid- to high-end FortiGate devices, including the 100D, 600C, 1000C, and 1240B.

To enable and configure the virtual switch - CLI:

```
config system virtual-switch
  edit vs1
    set physical-switch sw0
  config port
    edit 1
      set port port1
      set speed xx
      set duplex xx
      set status [up|down]
    edit 2
      set port port2
      set ...
    next
  next
end
```

Support for 802.1x fallback and 802.1x dynamic VLANs

There are four modes when enabling 802.1x on a virtual switch interface:

Mode	Description
Default	In this mode, it works as it did previously.
Fallback	In fallback mode, the virtual switch is treated as a master. Only one slave can refer to a fallback master. The ports in the master virtual switch are always authorized. After passing 802.1x authentication, the ports will stay authorized and move to its slave virtual switch.
Dynamic-vlan	In dynamic-vlan mode, the virtual switch is also treated as a master. However, many slaves can refer to a dynamic-vlan master. Those ports in the master virtual switch are always unauthorized. After passing 802.1x/MAB authentication, the ports are set to authorized and moved to one of its slave virtual switches.
Slave	In slave mode, a master must be set through security-8021x-master attribute. A slave virtual switch will use its master virtual switch's security-groups settings for authentication.

CLI example for fallback mode:

```

config system virtual-switch
  edit "fallsw"
    set physical-switch "sw0"
    config port
  next
  edit "trust"
    set physical-switch "sw0"
  next
config system interface
  edit "fallsw"
    set vdom "root"
    set ip 192.168.20.1 255.255.255.0
    set allowaccess ping https ssh snmp http telnet fgfm auto-ipsec radius-acct
    probereponse capwap
    set type hard-switch
    set security-mode 802.1X
    set security-8021x-mode fallback(fallback mode master switch)
    set security-groups "rds-grp" (the usergroup for 802.1x)
    set snmp-index 10
  next
  edit "trust"
    set vdom "root"
    set ip 192.168.22.1 255.255.255.0
    set allowaccess ping https ssh snmp http telnet fgfm auto-ipsec radius-acct
    probereponse
    set type hard-switch
    set security-mode 802.1X
    set security-8021x-mode slave(slave mode switch)
    set security-8021x-master "fallsw" (its master switch)
    set snmp-index 6
  next
end

```

CLI example for dynamic-vlan mode:

```
config system virtual-switch
edit "internal"
    set physical-switch "sw0"
edit "lan-trust"
    set physical-switch "sw0"
next
edit "lan-vlan1000"
    set physical-switch "sw0"
next
edit "lan-vlan2000"
    set physical-switch "sw0"
    config port
    edit "internal1" (normally we should not add port in slave switch. This is used if
        user wants to manually add one port in slave)
    next
end
config system interface
edit "internal"
    set vdom "root"
    set ip 192.168.11.99 255.255.255.0
    set allowaccess ping https ssh http fgfm capwap
    set type hard-switch
    set security-mode 802.1X
    set security-8021x-mode dynamic-vlan<-----dynamic-vlan mode master switch
    set security-groups "rds-grp"<-----the usergroup for 802.1x
    set snmp-index 15
next
edit "lan-trust"
    set vdom "root"
    set ip 192.168.111.99 255.255.255.0
    set allowaccess ping https ssh snmp http telnet fgfm auto-ipsec radius-acct
        proberesponse capwap
    set type hard-switch
    set security-mode 802.1X
    set security-8021x-mode slave<-----slave mode switch
    set security-8021x-master "internal"<-----its master switch
    set snmp-index 7
next
edit "lan-vlan1000"
    set vdom "root"
    set ip 192.168.110.1 255.255.255.0
    set allowaccess ping https ssh snmp http telnet fgfm auto-ipsec radius-acct
        proberesponse capwap
    set type hard-switch
    set security-mode 802.1X
    set security-8021x-mode slave<-----slave mode switch
    set security-8021x-master "internal"<-----its master switch
    set security-8021x-dynamic-vlan-id 1000 <-----the matching vlan id for this virtual
        switch
    set snmp-index 16
next
edit "lan-vlan2000"
    set vdom "root"
    set ip 192.168.220.1 255.255.255.0
```

```

        set allowaccess ping https ssh snmp http telnet fgfm auto-ipsec radius-acct
        proberesponse
    capwap
    set type hard-switch
    set security-mode 802.1X
    set security-8021x-mode slave
    set security-8021x-master "internal"
    set security-8021x-dynamic-vlan-id 2000
    set snmp-index 17
end
config user group
    edit "rds-grp"
        set dynamic-vlan-id 4000 (default vlan id if there is no vlan attribute return
        from server)
        set member "190"
    end
end

```

Zones

Zones are a group of one or more FortiGate interfaces, both physical and virtual, that you can apply security policies to to control inbound and outbound traffic. Grouping interfaces and VLAN subinterfaces into zones simplifies the creation of security policies where a number of network segments can use the same policy settings and protection profiles. When you add a zone, you select the names of the interfaces and VLAN subinterfaces to add to the zone. Each interface still has its own address and routing is still done between interfaces, that is, routing isn't affected by zones. You can create security policies to control the flow of intra-zone traffic.

For example, the network includes three separate groups of users representing different entities on the company network. While each group has its own set of port and VLANs, in each area, they can all use the same security policy and protection profiles to access the Internet. Rather than an administrator making nine separate security policies, an administrator can add the required interfaces to a zone and create three policies, making administration simpler.

You can configure policies for connections to and from a zone, but not between interfaces in a zone.

The following example shows how to set up a zone to include the internal interface and a VLAN.

To create a zone - GUI:

1. Go to **Network > Interfaces**.
2. Select the arrow on the **Create New** button and select **Zone**.
3. Enter a zone name of `Zone_1`.
4. Select the required **Interface Members**.
5. Select **OK**.

To create a zone - CLI:

```

config system zone
    edit Zone_1
        set interface internal VLAN_1
    next
end

```


Virtual domains

Virtual domains (VDMs) are a method of dividing a FortiGate into two or more virtual units that function as multiple independent units. A single FortiGate is then flexible enough to serve multiple departments of an organization, separate organizations, or to act as the basis for a service provider's managed security service.

VDMs provide separate security domains that allow separate zones, user authentication, security policies, routing, and VPN configurations. By default, each FortiGate has a VDM named root. This VDM includes all of the FortiGate physical interfaces, modem, virtual LAN (VLAN) subinterfaces, zones, security policies, routing settings, and VPN settings.

When a packet enters a VDM, it's confined to that VDM. In a VDM, you can create security policies for connections between VLAN subinterfaces or zones in the VDM. Packets don't cross the virtual domain border internally. To travel between VDMs, a packet must pass through a firewall on a physical interface. The packet then arrives at another VDM on a different interface, but it must pass through another firewall before entering the VDM. Both VDMs are on the same FortiGate. Inter-VDMs change this behavior because they are internal interfaces. However, their packets go through all the same security measures as on physical interfaces.

The following example shows how to enable VDMs on a FortiGate and the basic and create a VDM accounting on the DMZ2 port and assign an administrator to maintain the VDM. First, enable VDMs on the FortiGate. When you enable VDMs, the FortiGate will log you out.

For desktop and low-end FortiGate devices, you use the CLI to enable VDMs. Once you enable VDMs, all further configuration can be done using the GUI or the CLI. On larger FortiGate units, you can use the GUI or the CLI to enable VDMs.

To enable VDMs - GUI:

1. Go to **Dashboard**.
2. In the **System Information** widget, select **Enable** for **Virtual Domain**.

The FortiGate logs you out. Once you log back in, you'll notice that the menu structure has changed. This reflects the global settings for all VDMs:

To enable VDMs - CLI:

```
config system global
    set vdom-admin enable
end
```

Next, add the VDM called accounting.

To add a VDM - GUI:

1. Go to **System > VDM**, and select **Create New**.
2. Enter the VDM name `accounting`.
3. Select **OK**.

To add a VDM - CLI:

```
config vdom
    edit <new_vdom_name>
end
```

With the VDOM created, you can assign a physical interface to it and assign it an IP address.

To assign physical interface to the accounting VDOM - GUI:

1. Go to **Network > Interfaces**.
2. Select the DMZ2 port row and select **Edit**.
3. For the **Virtual Domain** drop-down list, select **accounting**.
4. Select the **Addressing Mode** of **Manual**.
5. Enter the IP address for the port of 10.13.101.100/24.
6. Set the **Administrative Access** to **HTTPS** and **SSH**.
7. Select **OK**.

To assign physical interface to the accounting VDOM - CLI:

```
config global
  config system interface
    edit dmz2
      set vdom accounting
      set ip 10.13.101.100/24
      set allowaccess https ssh
    next
  end
```

Wireless

A wireless interface is similar to a physical interface except it doesn't include a physical connection. FortiWiFi devices allow you to add multiple wireless interfaces that can be available at the same time. On FortiWiFi units, you can configure the device to be either an access point or a wireless client. As an access point, the FortiWiFi unit can have separate SSIDs, each on their own subnet for wireless access. In client mode, the FortiWiFi has only one SSID, and is used as a receiver to allow remote users to connect to the existing network using wireless protocols.

Wireless interfaces also require additional security measures to ensure the session doesn't get hijacked and data tampered with or stolen.

For more information about configuring wireless interfaces see the [FortiOS FortiWiFi and FortiAP Configuration Guide](#).

VLANs

Virtual Local Area Networks (VLANs) multiply the capabilities of a FortiGate, and can also provide added network security. VLANs use ID tags to logically separate devices on a network into smaller broadcast domains. These smaller domains forward packets only to devices that are part of that VLAN domain. This reduces traffic and increases network security. The IEEE 802.1Q standard defines VLANs. All layer-2 and layer-3 devices along a route must be 802.1Q-compliant to support VLANs along that route.

A Local Area Network (LAN) is a group of connected computers and devices that are arranged into network broadcast domains. A LAN broadcast domain includes all the computers that receive a packet broadcast from any computer in that broadcast domain. A switch automatically forwards the packets to all of its ports. In contrast, routers don't automatically forward network broadcast packets. This means routers separate broadcast domains. If a network has only switches and no routers, that network is considered one broadcast domain, no matter how

large or small it is. Smaller broadcast domains are more efficient because fewer devices receive unnecessary packets. They are more secure as well because a hacker reading traffic on the network will have access to only a small portion of the network instead of the entire network's traffic.

VLANs reduce the size of the broadcast domains by only forwarding packets to interfaces that are part of that VLAN or part of a VLAN trunk link. Trunk links form switch-to-switch or switch-to-router connections, and forward traffic for all VLANs. This enables a VLAN to include devices that are part of the same broadcast domain, but physically distant from each other.

VLAN ID tags consist of a 4-byte frame extension that switches and routers apply to every packet sent and received in the VLAN. Workstations and desktop computers, which are commonly originators or destinations of network traffic, aren't an active part of the VLAN process. All of the VLAN tagging and tag removal is done after the packet has left the computer.

A FortiGate that doesn't have VDOMs enabled can have a maximum of 255 interfaces in transparent operating mode. The same is true for any single VDOM. In NAT mode, the number can range from 255 to 8192 interfaces per VDOM, depending on the FortiGate model. These numbers include VLANs, other virtual interfaces, and physical interfaces. To have more than 255 interfaces configured in transparent operating mode, you need to configure multiple VDOMs that enable you to divide the total number of interfaces over all the VDOMs.

One example of an application of VLANs is a company's accounting department. Accounting computers may be located at both main and branch offices. However, accounting computers need to communicate with each other frequently and require increased security. VLANs allow the accounting network traffic to be sent only to accounting computers and to connect accounting computers in different locations as if they were on the same physical subnet.

This guide uses the term "packet" to refer to both layer-2 frames and layer-3 packets.

On a layer-2 switch, you can have only one VLAN subinterface per physical interface, unless that interface is configured as a trunk link. Trunk links can transport traffic for multiple VLANs to other parts of the network.

You can add multiple VLANs to the same physical interface on a FortiGate. However, VLAN subinterfaces added to the same physical interface can't have the same VLAN ID or have IP addresses on the same subnet. You can add VLAN subinterfaces with the same VLAN ID to different physical interfaces.

Creating VLAN subinterfaces with the same VLAN ID doesn't create an internal connection between them. For example, a VLAN ID of 300 on port1 and VLAN ID of 300 on port2 are allowed, but they aren't connected. Their relationship is the same as between any two FortiGate network interfaces.

FortiGate interfaces can't have overlapping IP addresses, the IP addresses of all interfaces must be on different subnets. This rule applies to both physical interfaces and to virtual interfaces, such as VLAN subinterfaces. Each VLAN subinterface must be configured with its own IP address and netmask. This rule helps prevent a broadcast storm or other similar network problems.

The following example shows how to add a VLAN, called `vlan_accounting`, on the FortiGate internal interface with an IP address of 10.13.101.101.

To add a VLAN - GUI:

1. Go to **Network > Interfaces**.
2. Select **Create New** and click on **Interface**.
The **Type** is set to VLAN, by default.
3. Enter a name for the **VLAN** to `vlan_accounting`.
4. Select the **Internal** interface.

5. Enter the VLAN ID.

The VLAN ID is a number between 1 and 4094 that allow groups of IP addresses with the same VLAN ID to be associated together.

6. Select the Addressing Mode of Manual.**7. Enter the IP address for the port of 10.13.101.101/24.****8. Set the Administrative Access to HTTPS and SSH.****9. Select OK.****To add a VLAN - CLI:**

```
config system interface
edit VLAN_1
set interface internal
set type vlan
set vlanid 100
set ip 10.13.101.101/24
set allowaccess https ssh
next
end
```

VLANs in NAT mode

In NAT mode, the FortiGate functions as a layer-3 device. In this mode, the FortiGate controls the flow of packets between VLANs, but can also remove VLAN tags from incoming VLAN packets. The FortiGate unit can also forward untagged packets to other networks, such as the Internet.

In NAT mode, the FortiGate supports VLAN trunk links with IEEE 802.1Q-compliant switches, or routers. The trunk link transports VLAN-tagged packets between physical subnets or networks. When you add VLAN sub-interfaces to the FortiGate physical interfaces, the VLANs have IDs that match the VLAN IDs of packets on the trunk link. The FortiGate directs packets with VLAN IDs to sub-interfaces with matching IDs.

You can define VLAN sub-interfaces on all FortiGate physical interfaces. However, if multiple virtual domains are configured on the FortiGate, you will have access to only the physical interfaces on your virtual domain. The FortiGate can tag packets leaving on a VLAN subinterface. It can also remove VLAN tags from incoming packets and add a different VLAN tag to outgoing packets.

Normally in VLAN configurations, the FortiGate device's internal interface is connected to a VLAN trunk, and the external interface connects to an Internet router that isn't configured for VLANs. In this configuration, the FortiGate can apply different policies for traffic on each VLAN interface connected to the internal interface, which results in less network traffic and better security.

Adding VLAN subinterfaces

A VLAN subinterface, also called a VLAN, is a virtual interface on a physical interface. The subinterface allows routing of VLAN tagged packets using that physical interface, but it's separate from any other traffic on the physical interface.

Adding a VLAN subinterface includes configuring:

- a physical interface
- an IP address and netmask
- a VLAN ID
- a VDOM

Physical interface

The term VLAN subinterface correctly implies the VLAN interface isn't a complete interface by itself. You add a VLAN subinterface to the physical interface that receives VLAN-tagged packets. The physical interface can belong to a different VDOM than the VLAN, but it must be connected to a network router that is configured for this VLAN. Without that router, the VLAN won't be connected to the network, and VLAN traffic won't be able to access this interface. The traffic on the VLAN is separate from any other traffic on the physical interface.

When you are working with interfaces on a FortiGate, use the **Column Settings** on the Interface display to make sure the information you need is displayed. When working with VLANs, it's useful to position the **VLAN ID** column close to the IP address. If you're working with VDOMs, including the **Virtual Domain** column as well will help you troubleshoot problems more quickly.

To view the Interface display, go to **Network > Interfaces**.

IP address and netmask

FortiGate interfaces can't have overlapping IP addresses. The IP addresses of all interfaces must be on different subnets. This rule applies to both physical and virtual interfaces, such as VLAN subinterfaces. Each VLAN subinterface must be configured with its own IP address and netmask pair. This rule helps prevent a broadcast storm or other similar network problems.



If you're unable to change your existing configurations to prevent IP overlap, enter the CLI command `config system settings and set allow-subnet-overlap enable` to allow IP address overlap. If you enter this command, multiple VLAN interfaces can have an IP address that's part of a subnet used by another interface. This command is recommended for advanced users only.

VLAN ID

The VLAN ID is part of the VLAN tag added to the packets by VLAN switches and routers. The VLAN ID is a number between 1 and 4094 that allow groups of IP addresses with the same VLAN ID to be associated together. VLAN ID 0 is used only for high priority frames, and 4095 is reserved.

All devices along a route must support the VLAN ID of the traffic along that route. Otherwise, the traffic will be discarded before reaching its destination. For example, if your computer is part of VLAN_100 and a co-worker on a different floor of your building is also on the same VLAN_100, you can communicate with each other over VLAN_100, only if all the switches and routers support VLANs and are configured to pass along VLAN_100 traffic properly. Otherwise, any traffic you send to your co-worker will be blocked or won't be delivered.

VDOM

If VDOMs are enabled, each VLAN subinterface must belong to a VDOM. This rule also applies for physical interfaces.



Interface-related CLI commands require a VDOM to be specified, regardless of whether a FortiGate has VDOMs enabled.

VLAN subinterfaces on separate VDOMs can't communicate directly with each other. In this situation, the VLAN traffic must exit the FortiGate and re-enter the unit, passing through firewalls in both directions. This situation is the same for physical interfaces.

A VLAN subinterface can belong to a different VDOM than the physical interface it is part of. This is because the traffic on the VLAN is handled separately from the other traffic on that interface. This is one of the main strengths of VLANs.

The following procedure will add a VLAN subinterface called `VLAN_100` to the FortiGate internal interface with a VLAN ID of 100. It will have an IP address and netmask of `172.100.1.1/255.255.255.0`, and allow HTTPS and PING administrative access. Note that in the CLI, you must enter “`set type vlan`” before setting the `vlanid`, and that the `allowaccess` protocols are lower case.

To add a VLAN subinterface in NAT mode - GUI:

1. If **Current VDOM** appears at the bottom left of the screen, select **Global** from the list of VDOMs.
2. Go to **Network > Interfaces**.
3. Select **Create New** to add a VLAN subinterface.
4. Enter the following:

VLAN Name	VLAN_100
Type	VLAN
Interface	internal
VLAN ID	100
Addressing Mod	Manual
IP/Netmask	172.100.1.1/255.255.255.0
Administrative Access	HTTPS, PING, TELNET

5. Select **OK**.

To view the new VLAN subinterface, select the expand arrow next to the parent physical interface (the internal interface). This will expand the display to show all VLAN subinterfaces on this physical interface. If there is no expand arrow displayed, there are no subinterfaces configured on that physical interface.

For each VLAN, the list displays the name of the VLAN, and, depending on column settings, its IP address, the Administrative access you selected for it, the VLAN ID number, and which VDOM it belongs to if VDOMs are enabled.

To add a VLAN subinterface in NAT mode - CLI:

```
config system interface
  edit VLAN_100
    set interface internal
    set type vlan
    set vlanid 100
    set ip 172.100.1.1 255.255.255.0
    set allowaccess https ping
  next
end
```

Configuring security policies and routing

Once you create a VLAN subinterface on a FortiGate, you need to configure security policies and routing for that VLAN. Without these, the FortiGate won't pass VLAN traffic to its intended destination. Security policies direct traffic through the FortiGate between interfaces. Routing directs traffic across the network.

Configuring security policies

Security policies permit communication between the FortiGate device's network interfaces, based on source and destination IP addresses. Interfaces that communicate with the VLAN interface need security policies to permit traffic to pass between them and the VLAN interface.

Each VLAN needs a security policy for each of the following connections the VLAN will be using:

- From this VLAN to an external network
- From an external network to this VLAN
- From this VLAN to another VLAN in the same virtual domain on the FortiGate
- From another VLAN to this VLAN in the same virtual domain on the FortiGate

The packets on each VLAN are subject to antivirus scans and other security profiles measures as they pass through the FortiGate.

Configuring routing

As a minimum, you need to configure a default static route to a gateway with access to an external network for outbound packets. In more complex cases, you must configure different static or dynamic routes based on packet source and destination addresses.

As with firewalls, you must configure routes for VLAN traffic. VLANs need routing and a gateway configured to send and receive packets outside their local subnet just as physical interfaces do. The type of routing you configure, static or dynamic, will depend on the routing used by the subnet and interfaces you're connecting to. Dynamic routing can be routing information protocol (RIP), border gateway protocol (BGP), open shortest path first (OSPF), or multicast.

If you enable SSH, PING, HTTPS and HTTP on the VLAN, you can use those protocols to troubleshoot your routing and test that it's properly configured. Enabling logging on the interfaces and using CLI diagnose commands, such as `diagnose sniff packet <interface_name>`, can also help locate any possible configuration or hardware issues.

VLANs in transparent mode

In transparent mode, a FortiGate behaves like a layer-2 bridge but it can still provide services such as antivirus scanning, web filtering, spam filtering, and intrusion protection to traffic. There are some limitations in transparent mode because you can't use SSL VPN, PPTP/L2TP VPN, DHCP server, or easily perform NAT on traffic. The limits in transparent mode apply to IEEE 802.1Q VLAN trunks passing through the device.

VLANs and transparent mode

You can insert a FortiGate operating in transparent mode into the VLAN trunk without making changes to your network. In a typical configuration, the FortiGate internal interface accepts VLAN packets on a VLAN trunk from a VLAN switch or router connected to internal network VLANs. The FortiGate external interface forwards VLAN-tagged packets through another VLAN trunk to an external VLAN switch or router and on to external

networks, such as the Internet. You can configure the unit to apply different policies for traffic on each VLAN in the trunk.

To pass VLAN traffic through the FortiGate, add two VLAN subinterfaces with the same VLAN ID, one to the internal interface and the other to the external interface. You then create a security policy to permit packets to flow from the internal VLAN interface to the external VLAN interface. If required, you create another security policy to permit packets to flow from the external VLAN interface to the internal VLAN interface. Typically in transparent mode, you don't permit packets to move between different VLANs. Network protection features, such as spam filtering, web filtering, and anti-virus scanning, are applied through the Security Profiles specified in each security policy, enabling very detailed control over traffic.

When the FortiGate receives a VLAN-tagged packet at a physical interface, it directs the packet to the VLAN subinterface with the matching VLAN ID. The VLAN tag is removed from the packet, and the FortiGate then applies security policies using the same method it uses for non-VLAN packets. If the packet exits the FortiGate through a VLAN subinterface, the VLAN ID for that subinterface is added to the packet and the packet is sent to the corresponding physical interface.

General configuration steps

There are two essential steps to configure a FortiGate to work with VLANs in transparent mode: add VLAN subinterfaces and create security policies.

You can also configure the Security Profiles that manage antivirus scanning, web filtering and spam filtering. For more information about Security profiles, see the [FortiOS Security Profiles Handbook](#).

Add VLAN subinterfaces

The VLAN ID of each VLAN subinterface must match the VLAN ID added by the IEEE 802.1Q-compliant router or switch. The VLAN ID can be any number between 1 and 4094, with 0 being used only for high priority frames and 4095 being reserved. You add VLAN subinterfaces to the physical interface that receives VLAN-tagged packets.

For this example, we're creating a VLAN called `internal_v225` on the internal interface, with a VLAN ID of 225. Administrative access is enabled for HTTPS and SSH. VDOMs aren't enabled.

To add VLAN subinterfaces in transparent mode - GUI:

1. Go to **Network > Interfaces**.
2. Select **Create New** and click on **Interfaces**.
3. Enter the following information and select **OK**.

Name	internal_v225
Type	VLAN
Interface	internal
VLAN ID	225
Administrative Access	Enable HTTPS, and SSH. These are very secure access methods.

The FortiGate adds the new subinterface to the interface that you selected.

Repeat steps 2 and 3 to add additional VLANs. You will need to change the **VLAN ID**, **Name**, and possibly **Interface** when adding additional VLANs.

To add VLAN subinterfaces in transparent mode - CLI:

```

config system interface
  edit internal_v225
    set interface internal
    set vlanid 225
    set allowaccess HTTPS SSH
    set description "VLAN 225 on internal interface"
    set vdom root
  end

```

Create security policies

In transparent mode, the FortiGate performs antivirus and antispam scanning on each VLAN's packets as they pass through the device. You need security policies to permit packets to pass from the VLAN interface where they enter the device to the VLAN interface where they exit the device. If there are no security policies configured, no packets will be allowed to pass from one interface to another.

To add security policies for VLAN subinterfaces - GUI:

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New** to add firewall addresses that match the source and destination IP addresses of VLAN packets.
3. Go to **Policy & Objects > IPv4 Policy** or **Policy & Objects > IPv6 Policy** and select **Create New**.
4. From the **Incoming Interface/Zone** list, select the VLAN interface where packets enter the unit.
5. From the **Outgoing Interface/Zone** list, select the VLAN interface where packets exit the unit.
6. Select the **Source** and **Destination** Address names that you added in step 2.
7. Select **OK**.

To add security policies for VLAN subinterfaces - CLI:

```

config firewall address
  edit incoming_VLAN_address
    set associated-interface <incoming_VLAN_interface>
    set type ipmask
    set subnet <IPv4_address_mask>
  next
  edit outgoing_VLAN_address
    set associated-interface <outgoing_VLAN_interface>
    set type ipmask
    set subnet <IPv4_address_mask>
  next
end
config firewall policy or config firewall policy6
  edit <unused_policy_number>
    set srcintf <incoming_VLAN_interface>
    set srcaddr incoming_VLAN_address
    set destintf <outgoing_VLAN_interface>
    set destaddr outgoing_VLAN_address
    set schedule always
    set service <protocol_to_allow_on_VLAN>
    set action ACCEPT
  next
end

```

VLANs over VXLANs

You can create a Virtual LAN (VLAN) over a Virtual Extensible LAN (VXLAN) tunnel interface. This allows you to keep VLAN tags in VXLAN traffic.

VLAN switching and routing

VLAN switching takes place on the Open Systems Interconnect (OSI) model layer-2, just like other network switching. VLAN routing takes place on the OSI model layer-3. The difference between them is that during VLAN switching, VLAN packets are simply forwarded to their destination. This is different from VLAN routing where devices can open the VLAN packets and change their VLAN ID tags to route the packets to a new destination.

VLAN layer-2 switching

Ethernet switches are layer-2 devices, and generally are 802.1Q compliant. Layer 2 refers to the second layer of the seven layer OSI basic networking model, called the Data Link layer. FortiGate devices act as layer-2 switches or bridges when they are in transparent mode. The devices simply tag and forward the VLAN traffic or receive and remove the tags from the packets. A layer-2 device doesn't inspect incoming packets or change their contents; it only adds or removes tags and routes the packet.

A VLAN can have any number of physical interfaces assigned to it. Multiple VLANs can be assigned to the same physical interface. Typically two or more physical interfaces are assigned to a VLAN, one for incoming and one for outgoing traffic. Multiple VLANs can be configured on one FortiGate, including trunk links.

VLAN layer-3 routing

Routers are layer-3 devices. Layer 3 refers to the third layer of the OSI networking model, the Network layer. FortiGate devices in NAT mode act as layer-3 devices. As with layer 2, FortiGate devices acting as layer-3 devices are 802.1Q-compliant.

The main difference between layer-2 and layer-3 devices is how they process VLAN tags. Layer-2 switches just add, read, and remove the tags. They don't alter the tags or do any other high-level actions. Layer-3 routers not only add, read, and remove tags but also analyze the data frame and its contents. This analysis allows layer-3 routers to change the VLAN tag if it's appropriate and send the data frame out on a different VLAN.

In a layer-3 environment, the 802.1Q-compliant router receives the data frame and assigns a VLAN ID. The router then forwards the data frame to other members of the same VLAN broadcast domain. The broadcast domain can include local ports, layer-2 devices and layer-3 devices, such as routers and firewalls. When a layer-3 device receives the data frame, the device removes the VLAN tag and examines its contents to decide what to do with the data frame. The layer-3 device considers:

- Source and destination addresses
- Protocol
- Port number

The data frame may be forwarded to another VLAN, sent to a regular non-VLAN-tagged network or just forwarded to the same VLAN as a layer-2 switch would do. Or, the data frame may be discarded if the proper security policy has been configured to do so.

Layer-2 and ARP traffic

By default, FortiGate devices don't pass layer-2 traffic. If there are layer-2 protocols such as IPX, PPTP or L2TP in use on your network, you need to configure the FortiGate interfaces to pass these protocols without blocking. Another type of layer-2 traffic is Address Resolution Protocol (ARP) traffic.

To allow layer 2 protocols - CLI:

```
config system interface
  edit <name_str>
    set l2forward enable
  end
```

where <name_str> is the name of an interface.

If VDOMs are enabled, this command is per VDOM. You must set it for each VDOM that has the problem as follows:

```
config vdom
  edit <vdom_name>
    config system interface
      edit <name_str>
        set l2forward enable
      end
    end
  end
```

If you enable layer-2 traffic, you may experience a problem if packets are allowed to repeatedly loop through the network. This repeated looping, very similar to a broadcast storm, occurs when you have more than one layer-2 path to a destination. Traffic may overflow and bring your network to a halt. You can break the loop by enabling Spanning Tree Protocol (STP) on your network's switches and routers.

STP forwarding

A FortiGate doesn't participate in the Spanning Tree Protocol (STP). STP is an IEEE 802.1 protocol that ensures there are no layer-2 loops on the network. Loops are created when there is more than one route for traffic to take and that traffic is broadcast back to the original switch. This loop floods the network with traffic, reducing available bandwidth to nothing.

If you use a FortiGate in a network topology that relies on STP for network loop protection, you need to make changes to the FortiGate configuration. Otherwise, STP recognizes the FortiGate as a blocked link and forwards the data to another path. By default, the FortiGate blocks STP as well as other non-IP protocol traffic.

Using the CLI, you can enable forwarding of STP and other layer-2 protocols through the interface. In this example, layer-2 forwarding is enabled on the external interface:

```
config system interface
  edit external
    set l2forward enable
    set stpforward enable
  next
end
```

By substituting different commands for `stpforward enable`, you can also allow layer-2 protocols, such as IPX, PPTP, or L2TP, to be used on the network.

STP support for FortiGate models with hardware switches

STP (Spanning Tree Protocol) used to be available only on the old style switch mode for the internal ports. You can now activate STP on the hardware switches found in the newer FortiGate models. These models use a virtual switch to simulate the old switch mode for the internal ports.

To enable STP - CLI:

```
config system interface
  edit lan
    set stp {enable | disable}
  next
end
```

ARP traffic

Address Resolution Protocol (ARP) packets are vital to communication on a network and ARP support is enabled on FortiGate interfaces, by default. Normally, you want ARP packets to pass through a FortiGate, especially if it's sitting between a client and a server or between a client and a router.

ARP traffic can cause problems, especially in transparent mode where ARP packets arriving on one interface are sent to all other interfaces including VLAN subinterfaces. Some layer-2 switches become unstable when they detect the same MAC address originating on more than one switch interface or from more than one VLAN. This instability can occur if the layer-2 switch doesn't maintain separate MAC address tables for each VLAN. Unstable switches may reset and cause network traffic to slow down considerably.

The default ARP timeout value is 5 minutes (300 seconds). ARP entries are usually removed after 5 minutes. However, some conditions can cause ARP entries to remain on the list for a longer period of time. This isn't a value that you can configure. To view the ARP list, enter the `get system arp` CLI command.

Proxy ARP extensions

You can extend the proxy ARP configuration to an IP address range instead of a single IP address. When you configure `proxy-arp`, in addition to setting the IP address, you can also set the `end-ip` address. If you don't set this, the proxy ARP will be a single address, as before. The following is an example CLI configuration, using the new setting:

```
config system proxy-arp
  edit 1
    set interface "internal"
    set ip 192.168.1.100
    set end-ip 192.168.1.102
  next
end
```

Multiple VDOMs solution

By default, physical interfaces are in the root domain. If you don't configure any of your VLANs in the root VDOM, it won't matter how many interfaces are in the root VDOM.

The multiple VDOMs solution is to configure multiple VDOMs on the FortiGate, one for each VLAN. In this solution, you configure one inbound and one outbound VLAN interface in each VDOM. ARP packets aren't forwarded between VDOMs. This configuration limits the VLANs in a VDOM and correspondingly reduces the administration needed per VDOM.

As a result of this configuration, the switches don't receive multiple ARP packets with duplicate MACs. Instead, the switches receive ARP packets with different VLAN IDs and different MACs. Your switches are stable.

However, you shouldn't use the multiple VDOMs solution under any of the following conditions:

- You have more VLANs than licensed VDOMs
- You don't have enough physical interfaces

Instead, use one of two possible solutions, depending on which operation mode you're using:

- In NAT mode, you can use the `vlanforward` CLI command.
- In transparent mode, you can use the `forward-domain` CLI command. But you still need to be careful in some rare configurations.

Vlanforward solution

If you're using NAT mode, the solution is to use the `vlanforward` CLI command for the interface in question. By default, this command is enabled and will forward VLAN traffic to all VLANs on this interface. When disabled, each VLAN on this physical interface can send traffic only to the same VLAN. There is no cross-talk between VLANs, and ARP packets are forced to take one path along the network which prevents the multiple paths problem.

In the following example, `vlanforward` is disabled on port1. All VLANs configured on port1 will be separate and will not forward any traffic to each other.

```
config system interface
  edit port1
    set vlanforward disable
  next
end
```

Forward-domain solution

If you're using transparent mode, the solution is to use the `forward-domain` CLI command. This command tags VLAN traffic as belonging to a particular collision group, and only VLANs tagged as part of that collision group receive that traffic. It's like an additional set of VLANs. By default, all interfaces and VLANs are part of forward-domain collision group 0. The many benefits of this solution include reduced administration, the need for fewer physical interfaces, and the availability of more flexible network solutions.

In the following example, forward-domain collision group 340 includes VLAN 340 traffic on port1 and untagged traffic on port 2. Forward-domain collision group 341 includes VLAN 341 traffic on port 1 and untagged traffic on port 3. All other interfaces are part of forward-domain collision group 0, by default. This configuration separates VLANs 340 and 341 from each other on port 1.

Use the following CLI commands:

```
config system interface
  edit port2
    set forward_domain 340
  next
  edit port3
    set forward_domain 341
  next
  edit port1-340
    set forward_domain 340
    set interface port1
```

```
        set vlanid 340
    next
    edit port1-341
        set forward_domain 341
        set interface port1
        set vlanid 341
    next
end
```

You may experience connection issues with layer-2 traffic, such as ping, if your network configuration has:

- Packets going through the FortiGate in transparent mode more than once
- More than one forwarding domain (such as incoming on one forwarding domain and outgoing on another)
- IPS and AV enabled

Now IPS and AV is applied the first time packets go through the FortiGate, but not on subsequent passes. Applying IPS and AV only to this first pass fixes the network layer-2 related connection issues.

Asymmetric routing

You might discover unexpectedly that hosts on some networks are unable to reach certain other networks. This occurs when request and response packets follow different paths. If a FortiGate recognizes the response packets, but not the requests, it blocks the packets as invalid. Also, if a FortiGate recognizes the same packets repeated on multiple interfaces, it blocks the session as a potential attack.

This is asymmetric routing. By default, a FortiGate blocks packets or drops the session when this happens. You can configure the FortiGate to permit asymmetric routing by using the following CLI commands:

```
config system settings
    set asymroute enable
end
```

If VDOMs are enabled, this command is per VDOM. You must set it for each VDOM that has the problem as follows:

```
config vdom
    edit <vdom_name>
        config system settings
            set asymroute enable
        end
    end
end
```

If this solves your blocked traffic issue, you know that asymmetric routing is the cause. But allowing asymmetric routing is not the best solution, because it reduces the security of your network.

For a long-term solution, it is better to change your routing configuration or change how the FortiGate connects to your network.



If you enable asymmetric routing, antivirus and intrusion prevention systems won't be effective. The FortiGate won't be aware of connections and will treat each packet individually. It will become a stateless firewall.

Configuring IPv4 and IPv6 ICMP traffic inspection

In order for the inspection of asymmetric ICMP traffic not to affect TCP and UDP traffic, you can enable or disable ICMP traffic inspection for traffic being routed asymmetrically for both IPv4 and IPv6.

To configure ICMP traffic inspection, use the following CLI commands:

- IPv4:

```
config system settings
  set asymroute-icmp
end
```

- IPv6:

```
config system settings
  set asymroute6-icmp
end
```

NetBIOS

Computers running Microsoft Windows operating systems that are connected through a network rely on a WINS server to resolve host names to IP addresses. The hosts communicate with the WINS server by using the NetBIOS protocol.

To support this type of network, you need to enable the forwarding of NetBIOS requests to a WINS server. The following example forwards NetBIOS requests on the internal interface for the WINS server located at an IP address of 192.168.111.222.

```
config system interface
  edit internal
    set netbios_forward enable
    set wins-ip 192.168.111.222
  next
end
```

These commands apply only to NAT mode. If VDOMs are enabled, these commands are per VDOM. You must set them for each VDOM that has the problem.

Too many VLAN interfaces

Any virtual domain can have a maximum of 255 interfaces in transparent mode. This includes VLANs, other virtual interfaces, and physical interfaces. NAT mode supports from 255 to 8192 depending on the FortiGate model. This total number of interfaces includes VLANs, other virtual interfaces, and physical interfaces.

A FortiGate may allow you to configure more interfaces than this. However, if you configure more than 255 interfaces, your system will become unstable and, over time, won't work properly. As all interfaces are used, they will overflow the routing table that stores the interface information, and connections will fail. When you try to add more interfaces, an error message will state that the maximum limit has already been reached.

If you see this error message, chances are you already have too many VLANs on your system and your routing has become unstable. To verify, delete a VLAN and try to add it back. If you have too many, you won't be able to add it back on to the system. In this case, you'll need to remove enough interfaces (including VLANs) so that the total number of interfaces drops to 255 or less. After doing this, you should also reboot the FortiGate to clean up its memory and buffers, or you will continue to experience unstable behavior.

To configure more than 255 interfaces on a FortiGate in transparent mode, you have to configure multiple VDOMs, each with many VLANs. However, if you want to create more than the default 10 VDOMs (or a maximum

of 2550 interfaces), you must buy a license for additional VDOMs and the FortiGate must be able to be licensed for more than 10 VDOMs.

With these extra licenses, you can configure up to 500 VDOMs, with each VDOM containing up to 255 VLANs in transparent mode. This is a theoretical maximum of over 127 500 interfaces. However, system resources will quickly get used up before reaching that theoretical maximum. To achieve the maximum number of VDOMs, you need to have top-end hardware with the most resources possible.

In NAT mode, if you have a top-end model, the maximum interfaces per VDOM can be as high as 8192, which is enough for all the VLANs in your configuration.



A FortiGate has limited resources, such as CPU load and memory, that are divided between all configured VDOMs. When running 250 or more VDOMs, you may need to monitor the system resources to ensure there is enough to support the configured traffic processing.

Troubleshooting VLAN issues

Several problems can occur with your VLANs. Since VLANs are interfaces with IP addresses, they behave as interfaces and can have similar problems that you can diagnose with tools, such as ping, traceroute, packet sniffing, and diag debug.

Enhanced MAC VLANs

The media access control (MAC) virtual local area network (VLAN) feature in Linux allows you to configure multiple virtual interfaces with different MAC addresses (and therefore different IP addresses) on a physical interface.

A FortiGate implements an enhanced MAC VLAN which consists of a MAC VLAN with bridge functionality. Because each MAC VLAN has a unique MAC address, virtual IP addresses (VIPs) and IP pools are supported, and you can disable Source Network Address Translation (SNAT) in policies. However, MAC VLAN can't be used in a transparent mode virtual domain (VDOM). In a transparent mode VDOM, a packet leaves an interface with the MAC address of the original source instead of the interface's MAC address. FortiGate solves this limitation by implementing an enhanced version of MAC VLAN, where it adds a MAC table in the MAC VLAN which learns various MAC addresses when traffic passes through.

If you configure a VLAN ID for an enhanced MAC VLAN, it won't join the switch of the underlying interface. When a packet is sent to this interface, a VLAN tag is inserted in the packet, and the packet is sent to the driver of the underlying interface. When the underlying interface receives a packet, if the VLAN ID doesn't match, it won't deliver the packet to this enhanced MAC VLAN interface.

If you use an interface in an enhanced MAC VLAN, you shouldn't use it for other purposes, such as a management interface, HA heartbeat interface, or in transparent VDOMs.

If a physical interface is used by an EMAC VLAN interface, you can't use it in a virtual wire pair.

To configure enhanced MAC VLAN – CLI:

```
config system interface
  edit <interface-name>
    set type emac-vlan
    set vlan-id <VLAN-ID>
    set interface <physical-interface>
```



```

next
end

```

Setting a VLAN ID is optional.

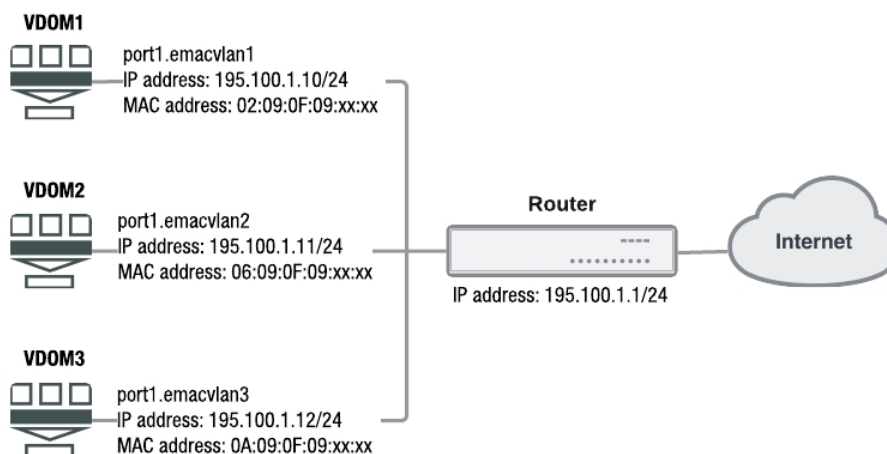
Enhanced MAC VLANs and HA

In high availability (HA) configurations, enhanced MAC VLAN is treated as a physical interface. It's assigned a unique physical interface ID and the MAC table is synchronized to the slaves in the same HA cluster.

Example 1: Enhanced MAC VLAN configuration for multiple VDOMs that use the same interface or VLAN

In this example, a FortiGate is connected, through port 1, to a router that's connected to the Internet. Three VDOMs share the same interface (port 1), which connects to the same router that's connected to the Internet. Three enhanced MAC VLAN interfaces are configured on port 1 for the three VDOMs. The enhanced MAC VLAN interfaces are in the same IP subnet segment and each have unique MAC addresses.

The underlying interface (port 1) can be a physical interface, an aggregate interface, or a VLAN interface on a physical or aggregate interface.



In this scenario, the configuration for enhanced MAC VLAN is the following:

```

config system interface
  edit port1.emacvlan1
    set vdom VDOM1
    set type emac-vlan
    set interface port1
  next
  edit port 1.emacvlan2
    set vdom VDOM2
    set type emac-vlan
    set interface port1
  next
  edit port1.emacvlan3
    set vdom VDOM3
    set type emac-vlan
    set interface port1
  next
end

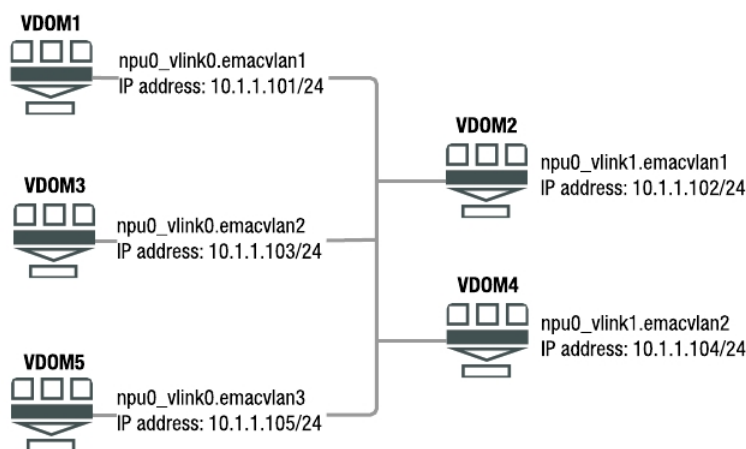
```

end

Example 2: Enhanced MAC VLAN configuration for shared VDOM links among multiple VDOMs

In this example, multiple VDOMs can connect to each other using enhanced MAC VLAN on network processing unit (NPU) virtual link (vlink) interfaces.

Currently, FortiGate VDOM links (npu-vlink) are designed to be peer-to-peer connections and VLAN interfaces on NPU vlink ports use the same MAC address. It's not practical to connect more than 2 VDOMs using NPU vlinks and VLAN interfaces.



In this scenario, the configuration for enhanced MAC VLAN is the following:

```

config system interface
  edit npu0_vlink0.emacvlan1
    set vdom VDOM1
    set type emac-vlan
    set interface npu0_vlink0
  next
  edit npu0_vlink0.emacvlan2
    set vdom VDOM3
    set type emac-vlan
    set interface npu0_vlink0
  next
  edit npu0_vlink1.emacvlan1
    set vdom VDOM2
    set type emac-vlan
    set interface npu0_vlink1
  next
end

```

Example 3: Enhanced MAC VLAN configuration for unique MAC addresses for each VLAN interface on the same physical port

In some networks, a unique MAC address is needed for each VLAN interface when the VLAN interfaces share the same physical port. In this case, the enhanced MAC VLAN interface is used the same way as normal VLAN

interfaces.

To configure this, you use the `set vlanid` command for the VLAN tag.

In this scenario, the configuration for enhanced MAC VLAN is the following:

```
config system interface
  edit interface-name
    set type emac-vlan
    set vlanid <VLAN-ID>
    set interface <physical-interface>
  end
```

Virtual wire pairs

A virtual wire pair logically binds two physical interfaces on a FortiGate, usually an internal and an external interface, together so that all traffic that one of the interfaces in a virtual wire pair accepts can exit the FortiGate only through the other interface in the virtual wire pair, and only if allowed by a virtual wire pair firewall policy. Traffic that arrives on other interfaces can't be routed to interfaces in a virtual wire pair.

The interfaces in a virtual wire pair don't have IP addresses, which means you can configure a virtual wire pair in your network without making any network changes. You can create more than one virtual wire pair on a FortiGate.

You can configure virtual wire pairs on a FortiGate that's running in either transparent or NAT modes. A virtual wire pair supports transparent mode between two interfaces without requiring you to change the FortiGate from NAT to Transparent mode.

If a physical interface is used by an EMAC VLAN interface, you can't use it in a virtual wire pair.

To configure a virtual wire pair - GUI:

Interfaces that you use for administrative access can't be used in a virtual wire pair. If you want to use an interface that you use for administrative access in a virtual wire pair, make sure you configure a different interface to allow administrative access before you create the virtual wire pair.

If the interfaces you want to use in a virtual wire pair are part of a switch, such as the default **lan** interface, you need to remove them from the switch before they can be added to the virtual wire pair.

1. Go to **Network > Interfaces** and select **Create New > Virtual Wire Pair**.
2. In the **Name** field, type a name for the virtual wire pair.
3. In the **Interface Members** field, select the interfaces that you want to add to the virtual wire pair.
4. If you want to enable wildcard VLANs for the virtual wire pair, enable **Wildcard VLAN**.
5. Select **OK**.

To configure a virtual wire pair policy - GUI:

1. Go **Policy & Objects > IPv4 Virtual Wire Pair Policy**.
2. Select a virtual wire pair in the upper right-hand corner of the screen, and select **Create New**.
3. In the **Name** field, type a name for the virtual wire pair policy.
4. In the **Virtual Wire Pair** field, select the direction that traffic is allowed to flow.
5. Configure other firewall options, as needed.
6. Select **OK**.
7. If necessary, create a second virtual wire pair policy to allow traffic to flow in the opposite direction.

Traffic can now flow through the FortiGate using the virtual wire pair. You can go to **FortiView > Policies** to see traffic flowing through both policies.



The **IPv4 Virtual Wire Pair Policy** menu item in the GUI appears only when you have created at least one virtual wire pair.

Wildcard VLANs for virtual wire pairs

Although you can't add virtual local area networks (VLANs) to virtual wire pairs, you can enable wildcard VLANs for a virtual wire pair. Doing this allows all VLAN-tagged traffic to pass through a virtual wire pair if a virtual wire pair firewall policy allows the traffic.

To enable wildcard VLANs for a virtual wire pair, enable the **Wildcard VLAN** option when you create a virtual wire pair.

VLAN filters for virtual wire pairs

After you enable wildcard VLANs, if you don't want a virtual wire pair policy to allow all VLAN traffic, you can specify VLAN filters. A VLAN filter allows only the VLANs in the filter and drops traffic with other VLAN tags. VLAN filters don't affect traffic that isn't VLAN-tagged. You configure VLAN filters using the CLI.

You can add a VLAN filter to a virtual wire pair to apply the filter to all traffic that the virtual wire pair accepts. You can also add a VLAN filter to a virtual wire pair firewall policy to apply more specific VLAN filtering only to the traffic that the policy accepts.

To configure VLAN filters for wildcard VLANs - CLI:

```
config system virtual-wire-pair
  edit <vwp-name>
    set member <vwp-interface1-name> <vwp-interface2-name>
    set wildcard-vlan enable
    set vlan-filter <VLAN-range-list>
  next
end

config firewall policy
  edit <policy-ID>
    set vlan-filter <VLAN-range-list>
  next
end
```



The `vlan-filter` option is only available for policies on virtual wire pairs that have the wildcard VLAN option enabled.

Botnet and command-and-control protection

You can configure botnet and command-and-control traffic protection, in a FortiGate GUI or CLI.

To configure botnet scans on an interface - GUI:

Select the **Scan Outgoing Connections to Botnet Sites** option on the **Interfaces** page. The options are **Disable**, **Block**, and **Monitor**.

To configure botnet scans on an interface - CLI:

```
config system interface
  edit <interface>
    set scan-botnet-connections {disable | block | monitor}
  next
end
```

You can also enable the scanning of botnet and command-and-control traffic in the following policies:

To enable botnet scans in firewall policies - CLI:

```
config firewall policy
  edit <policy ID>
    set scan-botnet-connections {disable | block | monitor}
  next
end
```

To enable botnet scans in firewall explicit proxy policies - CLI:

```
config firewall explicit-proxy-policy
  edit <policy ID>
    set scan-botnet-connections {disable | block | monitor}
  next
end
```

To enable botnet scans in firewall interface policies - CLI:

```
config firewall interface-policy
  edit <policy ID>
    set scan-botnet-connections {disable | block | monitor}
  next
end
```

To enable botnet scans for firewall sniffer - CLI:

```
config firewall sniffer
  edit <policy ID>
    set scan-botnet-connections {disable | block | monitor}
  next
end
```

DNS

A Domain Name System (DNS) server is a public service that converts symbolic node names to IP addresses. A DNS server implements the protocol. In simple terms, it acts as a phone book for the Internet. A DNS server matches domain names with their computer IP addresses. This allows you to use readable locations, such as `fortinet.com`, when you browse the Internet. FortiOS supports DNS configuration for both IPv4 and IPv6 addressing.

FortiGate includes default DNS server addresses. However, you should change these addresses to ones that your Internet Service Provider (ISP) provides. The defaults are DNS proxies and are not as reliable as those from your ISP.

Within FortiOS, there are two DNS configuration options. Each option provides a specific service and both options can work together to provide a complete DNS solution.

DNS settings

You configure basic DNS queries on interfaces that connect to the Internet. When a user requests a website, FortiGate looks to the configured DNS servers to provide the IP address of the website in order to know which server to contact to complete the transaction.

You configure DNS server addresses by selecting **Network > DNS**, and then specifying the DNS server addresses. These addresses are typically supplied by your ISP. If you have local Microsoft domains on the network, you can enter a domain name in the **Local Domain Name** field.

In a situation where all three fields are configured, FortiGate first looks to the local domain. If no match is found, FortiGate sends a request to the external DNS servers.

If virtual domains (VDOM) are enabled, you create a DNS database in each VDOM. All of the interfaces in a VDOM share the DNS database in that VDOM.

Additional DNS CLI configuration

Additional DNS configuration options are available in the CLI, using the `config system dns` command. Within this command, you can also set the following commands:

Command	Description
<code>dns-cache-limit</code>	Set how many DNS entries are stored in the cache. Entries that remain in the cache provide a quicker response to requests than going out to the Internet to get the same information.
<code>dns-cache-ttl</code>	Set how long entries remain in the cache, in seconds. Possible values are 60 to 86400 (default is 24 hours).
<code>cache-notfound-responses</code>	When you enable this, any DNS requests that are returned with NOT FOUND can be stored in the cache.
<code>source-ip</code>	Define a dedicated IP address for communications with the DNS server.

DDNS

If your ISP changes your external IP address on a regular basis, and you have a static domain name, you can configure the external interface to use a dynamic DNS (DDNS) service. This ensures that external users and customers can always connect to your company firewall. If you have a FortiGuard subscription, you can use FortiGuard as the DDNS server.

You can configure FortiGuard as the DDNS server, in the FortiGate GUI or CLI.

To configure FortiGuard as the DDNS server in the FortiGate GUI, select **Network > DNS** and enable **FortiGuard DDNS**. Then select the interface with the dynamic connection, which DDNS server you have an account with, your domain name, and account information. If your DDNS server isn't on the list, there is a generic option where you can provide your DDNS server information.

To configure FortiGuard as the DDNS server - CLI:

```
config system fortiguard
    set ddns-server-ip
    set ddns-server-port
end
```

If you don't have a FortiGuard subscription or want to use a different DDNS server, you can configure DDNS in the CLI. You can configure a DDNS for each interface. Only the first configured port appears in the FortiGate GUI. Additional commands vary depending on the DDNS server you select. Use the following CLI commands:

```
config system ddns
    edit <DDNS_ID>
        set monitor-interface <external_interface>
        set ddns-server <ddns_server_selection>
    next
end
```

Configuring FortiGate to refresh DDNS IP addresses

You can configure FortiGate to refresh DDNS IP addresses. FortiGate periodically checks the DDNS server that is configured.

To configure FortiGate to refresh DDNS IP addresses - CLI:

```
config system ddns
    edit <1>
        set ddns-server FortiGuardDDNS
        set use-public-ip enable
        set update-interval seconds
    next
end
```

The possible values for update-interval are 60 to 2592000 seconds, and the default is 300 seconds.

TLS support for DDNS updates

When cleartext is disabled, FortiGate uses the SSL connection to send and receive Dynamic DNS services (DDNS) updates.

To disable cleartext - CLI:

```
config system ddns
    set clear-text disable
end
```

You can also set the ssl-certificate name in the same location, using the following command:

```
set ssl-certificate <cert_name>
```

DDNS update override for DHCP

DHCP server has an override command option that allows DHCP server communications to go through DDNS to perform updates for the DHCP client. This enforces a DDNS update of the AA field every time, even if the DHCP client does not request it. This allows the support of the allow/ignore/deny client-updates options.

To enable DDNS update override - CLI:

```
config system dhcp server
    edit <0>
        set ddns-update_override enable
    next
end
```

FortiDDNS registration to a public IP address

Fortinet's Dynamic DNS services (FortiDDNS) can be registered to a public IP address even if the FortiGate model doesn't have any physical interfaces on the Internet. This applies to when the FortiGate is behind other networking devices that are employing NAT. You can configure this in the GUI and the CLI.

DNS servers

You can also create local DNS servers for your network. Depending on your requirements, you can manually maintain your entries (master DNS server) or use it as a jumping point, where the server refers to an outside source (slave DNS server). A local master DNS server works similarly to the DNS server addresses configured in **Network > DNS**, but you must manually add all entries. This allows you to add a local DNS server to include specific URL and IP address combinations.

The DNS server options are not visible in the FortiGate GUI, by default. To enable the server, select **System > Feature Visibility**, select **DNS Database**, and select **Apply**.

While a master DNS server is an easy method to include regularly used addresses to save on going to an outside DNS server, it isn't recommended to make it the authoritative DNS server. IP addresses may change and maintaining any type of list can become labor-intensive.

It's best to use a FortiGate master DNS server for local services. For example, a company has a web server in their DMZ that internal users (employees) and external users (customers or remote employees) access. When internal users access a website, a request for the website is sent out to the DNS server on the Internet, which then returns an IP address or virtual IP address. After the company configures an internal DNS server, the same website request is resolved internally to the internal web server IP address. This minimizes inbound and outbound traffic, and access time.

As a slave DNS server, a FortiGate refers to an external or alternate source as a way to obtain the URL and IP address combination. This is useful if there is a master DNS server for a large company, where a list is maintained. Satellite offices can then connect to the master DNS server to obtain the correct addressing.

The DNS server entries don't allow CNAME entries, as per [RFC 1912](#), section 2.4.

Configure a master DNS server - GUI:

1. Select **Network > DNS Servers**, and select **Create New** for **DNS Database**.
2. Select the **Type** of **Master**.
3. Select the **View** as **Shadow**.
4. The view is the accessibility of the DNS server. Selecting **Public**, external users can access, or use, the DNS server. Selecting **Shadow**, only internal users can use it.
5. Enter the **DNS Zone**, for example, `WebServer`.
6. Enter the domain name for the zone, for example `example.com`.
7. Enter the hostname of the DNS server, for example, `Corporate`.
8. Enter the contact address for the administrator, for example, `admin@example.com`.
9. Set **Authoritative** to **Disable**.
10. Select **OK**.
11. Enter the DNS entries for the server by selecting **Create New**.
12. Select the **Type**, for example, **Address (A)**.
13. Enter the **Hostname**, for example `web.example.com`.
14. Enter the remaining information, which varies depending on the **Type** selected.
15. Select **OK**.

Configure a master DNS server - CLI:

```
config system dns-database
  edit WebServer
    set domain example.com
    set type master
    set view shadow
    set ttl 86400
    set primary-name corporate
    set contact admin@example.com
    set authoritative disable
    config dns-entry
      edit 1
        set hostname web.example.com
        set type A
        set ip 192.168.21.12
        set status enable
      next
    next
  next
end
```

Configuring a recursive DNS

You can set an option to ensure this type of DNS server isn't the authoritative server. When configured, a FortiGate checks its internal DNS server (master or slave). If the request can't be fulfilled, it will look to the external DNS servers. This is known as a split DNS configuration.

You can also have a FortiGate look to an internal server if the master or slave doesn't fulfill the request, using the following CLI commands:

```
config system dns-database
  edit example.com
    ...
    set view shadow
  next
end
```

For this behavior to work completely, you must set the DNS query for the external interface to be recursive.

To configure a recursive DNS - GUI:

1. Go to **Network > DNS Servers**, and select **Create New** for **DNS Service on Interface**.
2. Select the **Interface**.
3. Select the **Mode** to **Recursive**.
4. Select **OK**.

To configure a recursive DNS - CLI:

```
config system dns-server
  edit wan1
    set mode recursive
  next
end
```

Configuring IPv6 Router Advertisement options for DNS configuration

FortiGate supports the following RFC 6106 IPv6 Router Advertisement options:

- Obtaining DNS search list options from upstream DHCPv6 servers
- Sending the DNS search list through Router Advertisement
- Sending the DNS search list through the FortiGate DHCP server
- Sending DNS search list option to downstream clients with Router Advertisements that use a static prefix (FortiOS version 5.6.1 and later)
- Sending recursive DNS server option to downstream clients with Router Advertisements that use a static prefix (FortiOS version 5.6.1 and later)

To obtain the DNS search list options from upstream DHCPv6 servers - CLI:

```
config system interface
  edit wan1
    config ipv6
      set dhcp6-prefix-delegation enable
    next
  next
end
```

```
end
```

To send DNS search lists through Router Advertisement - CLI:

```
config system interface
edit port 1
config IPv6
set ip6-address 2001:10::/64
set ip6-mode static
set ip6-send-adv enable
config ip6-delegated-prefix-list
edit 1
set upstream-interface WAN
set subnet 0:0:0:11::/64
set autonomous-flag enable
set onlink-flag enable
next
next
next
end
```

To send the DNS search lists through the FortiGate DHCP server - CLI:

You can use the `dns-search-list delegated` command to send DNS search list option to downstream clients with Router Advertisements that use a static prefix.

```
config system dhcp6 server
edit 1
set interface port2
set upstream-interface WAN
set ip-mode delegated
set dns-service delegated
set dns-search-list delegated
set subnet 0:0:0:12::/64
next
end
```

To send DNS search list option to downstream clients with Router Advertisements that use a static prefix - CLI:

You can use the `set dnssl <DNS search list option>` command to send DNS search list option to downstream clients with Router Advertisements that use a static prefix.

```
config system interface
edit port1
config ipv6
config ip6-prefix-list
edit <2001:db8::/64>
set autonomous-flag enable
set onlink-flag enable
set rdns 2001:1470:8000::66 2001:1470:8000::72
set dnssl <DNS search list option>
next
next
next
next
end
```

To send recursive DNS server option to downstream clients with Router Advertisements that use a static prefix - CLI:

You can use the `set rdns <recursive DNS search option>` command to send Recursive DNS server option to downstream clients with Router Advertisements that use a static prefix.

```
config system interface
  edit port1
    config ipv6
      config ip6-prefix-list
        edit <2001:db8::/64>
          set autonomous-flag enable
          set onlink-flag enable
          set rdns 2001:1470:8000::66 2001:1470:8000::72
          set dnssl <DNS search list option>
        next
      next
    next
  next
end
```

Internet services

The Internet Service Database (ISDB) is a database that contains a list of IP addresses, IP protocols, and port numbers that are used by the most common Internet services.

The IP Reputation Database (IRDB) is a database that's populated by the FortiGuard IP Reputation Service which aggregates malicious source IP data from the Fortinet distributed network of threat sensors and other global sources that collaborate to provide up-to-date threat intelligence about hostile sources. You can select categories, such as Proxy IP, Spam, and TOR Exit Node, to see specific information about each category.

A FortiGate regularly updates new versions of both the ISDB and IRDB from FortiGuard.

To view the ISDB - GUI:

1. Go to **Policy & Objects > Internet Service Database**.
2. In the **Name** column, expand **Internet Service Database**.

To view the IRDB - GUI:

1. Go to **Policy & Objects > Internet Service Database**.
2. In the **Name** column, expand **IP Reputation Database**.

You can use Internet services as a source in firewall policies. For more information, see the [FortiOS Firewall Handbook](#). You can also use Internet services as a source and destination in traffic shaping policies. For more information, see the [FortiOS Traffic Shaping Handbook](#).

Advanced static routing

Advanced static routing includes features and concepts that are used in more complex networks.

Routing concepts

Many routing concepts apply to static routing. However, without first understanding these basic concepts, it's difficult to understand the more complex dynamic routing.

Routing in VDOMs

Routing on FortiGate devices is configured per VDOM. This means if VDOMs are enabled, you must enter a VDOM to do any routing configuration. This allows each VDOM to operate independently, with its own default routes and routing configuration.

In this guide, the procedures assume that the FortiGate has VDOMs disabled. This is stated in the assumptions for the examples. If the FortiGate has VDOMs enabled, you'll need to perform the following steps in addition to the procedure steps.

To route in VDOMs - GUI:

Select the VDOM that you want to view or configure at the bottom of the main menu.

To route in VDOMs - CLI:

Before you follow any CLI routing procedures with VDOMs enabled, enter the following commands. For this example, it's assumed that you'll be working in the root VDOM. Change root to the name of your selected VDOM as needed.

```
config vdom
edit root
```

Following these commands, you can enter any routing CLI commands, as normal.

Default route

The default route is used if there are no other routes in the routing table or if none of the other routes apply to a destination. Including the gateway in the default route gives all traffic a next-hop address to use when leaving the local network. The gateway address is normally another router on the edge of the local network.

All routers, including FortiGate devices, are shipped with default routes in place. This allows you to set up and become operational more quickly. Beginner administrators can use the default route settings until a more advanced configuration is needed.

Adding or editing a static route

1. Go to **Network > Static Routes** and select **Create New**.
2. Enter the following information and select **OK**.

Destination IP/Mask	Enter the destination IP address and netmask. A value of 0 . 0 . 0 . 0 / 0 . 0 . 0 . 0 is universal.
Gateway	Enter the gateway IP address.
Interface	Select the name of the interface that the static route will connect through.
Administrative Distance	Enter the distance value, which will affect which routes are selected first by different protocols for route management or load balancing. The default is 10.
Advanced Options	Optionally, expand Advanced Options and enter a Priority , which will artificially weight the route during route selection. The higher the priority number, the less likely the route is to be selected over other routes. The default is 0.

Enabling or disabling individual static routes

You can enable or disable individual static routes.

To configure IPv4 static routes - CLI:

```
config router static
  edit <sequence number>
    set status {enable | disable}
  next
end
```

To configure IPv6 static routes - CLI:

```
config router static6
  edit <sequence number>
    set status {enable | disable}
  next
end
```

Configuring FQDNs as a destination address in static routes

You can configure FQDN firewall addresses as destination addresses in a static route, using either the GUI or the CLI.

In the GUI, to add an FQDN firewall address to a static route in the firewall address configuration, enable the **Static Route Configuration** option. Then, when you configure the static route, set **Destination** to **Named Address**.

In the CLI, use the following CLI commands:

First, configure the firewall FQDN address:

```
config firewall address
  edit 'Fortinet-Documentation-Website'
    set type fqdn
    set fqdn docs.fortinet.com
    set allow-routing enable
```

Next, add the FQDN address to a static route.

```
config router static
  edit 0
    set dstaddr Fortinet-Documentation-Website
    ...
  end
```

Routing table

When two computers are directly connected, there's no need for routing because each computer knows exactly where to find the other computer, and they communicate directly.

Networking computers allows many computers to communicate with each other. This requires each computer to have an IP address to identify its location to the other computers. This is much like a mailing address, where you won't receive your postal mail at home if you don't have an address for people to send mail to. The routing table on a computer is much like an address book used to mail letters to people, where the routing table maintains a list of how to reach computers. Routing tables may also include information about the quality of service (QoS) of the route, and the interface associated with the route if the device has multiple interfaces.

Looking at routing as delivering letters is more simple than reality. In reality, routers lose power or have bad cabling, network equipment is moved without warning, and other such events happen that prevent static routes from reaching their destinations. When any changes, such as these, happen along a static route, traffic can no longer reach the destination and the route goes down. Dynamic routing can address these changes to ensure that traffic still reaches its destination. The process of realizing there's a problem, backtracking, and finding a route that is operational, is called convergence. If there's fast convergence in a network, users won't even know that re-routing is taking place.

The routing table for any device on the network has a limited size. For this reason, routes that aren't used are replaced by new routes. This method ensures the routing table is always populated with the most current and most used routes, which are the routes that have the best chance of being reused. Another method that's used to maintain the routing table's size is if a route in the table and a new route are to the same destination, one of the routes is selected as the best route to that destination and the other route is discarded.

Routing tables are also used in unicast reverse path forwarding (uRPF). In uRPF, the router not only looks up the destination information but it also looks up the source information to ensure that it exists. If there's no source to be found, that packet is dropped because the router assumes it's an error or an attack on the network.

The routing table is used to store routes that are learned. The routing table for any device on the network has a limited size. For this reason, routes that aren't used are replaced by new routes. This method ensures the routing table is always populated with the most current and most used routes, which are the routes that have the best chance of being reused. Another method used to maintain the routing table's size is if a route in the table and a new route are to the same destination, one of the routes is selected as the best route to that destination and the other route is discarded.

Viewing the routing table

You can view the routing table in the FortiGate GUI. By default, all routes are displayed in the Routing Monitor list. The default static route is defined as 0.0.0.0/0, which matches the destination IP address of "any/all" packets.

To display the routes in the routing table, go to **Monitor > Routing Monitor**. Select **Static & Dynamic** to view the routes.

You can also monitor policy routes. Select **Policy** to list the active policy routes on the FortiGate and view information about them. The active policy routes include policy routes that you create, SD-WAN rules, and Internet service static routes. It also supports downstream devices in the Security Fabric.

The following figure show an example of the static and dynamic routes in the Routing Monitor list:

Refresh	Route Lookup	Edit Route	Create Address	Static & Dynamic Policy	
IP Version	Type	Network	Gateway IP	Interfaces	Distance
4	Static	0.0.0.0/0	172.25.176.1	port1	10
4	Connected	172.25.176.0/24	0.0.0.0	port1	0
6	Connected	::1/128	::	root	0
6	System	ff00::8	::	port3	0

The following figure show an example of the policy routes in the Routing Monitor list:

Field	Description
IP Version	Shows whether the route is IPv4 or IPv6. IPv6 routes are displayed only if IPv6 is enabled in the FortiGate GUI.
Type	<p>The type values assigned to FortiGate routes (Static, Connected, RIP, OSPF, or BGP).</p> <ul style="list-style-type: none"> • All: All routes recorded in the routing table • Connected: All routes associated with direct connections to FortiGate interfaces • Static: The static routes that have been added to the routing table manually • RIP: All routes learned through RIP. For more information, see "RIP" on page 2110. • RIPNG: All routes learned through RIP version 6 (which enables the sharing of routes through IPv6 networks) • BGP: All routes learned through BGP. For more information, see "BGP" on page 2187. • OSPF: All routes learned through OSPF. For more information, see "OSPF" on page 2146. • OSPF6: All routes learned through OSPF version 6 (which enables the sharing of routes through IPv6 networks) • IS-IS: All routes learned through IS-IS. For more information, see "IS-IS" on page 2224. • HA: RIP, OSPF, and BGP routes synchronized between the primary unit and the subordinate units of a high availability (HA) cluster. HA routes are maintained on subordinate units and are visible only if you're viewing the router monitor from a virtual domain that is configured as a subordinate virtual domain in a virtual cluster. <p>For more information about HA routing synchronization, see the FortiOS High Availability Handbook.</p>

Field	Description
Subtype	<p>If applicable, the subtype classification assigned to OSPF routes.</p> <p>An empty string implies an intra-area route. The destination is in an area that the FortiGate is connected to.</p> <ul style="list-style-type: none"> • OSPF inter area: The destination is in the OSPF AS, but FortiGate isn't connected to that area. • External 1: The destination is outside the OSPF AS. This is known as OSPF E1 type. The metric of a redistributed route is calculated by adding the external cost and the OSPF cost together. • External 2: The destination is outside the OSPF AS. This is known as OSPF E2 type. In this case, the metric of the redistributed route is equivalent to the external cost only, expressed as an OSPF cost. • OSPF NSSA 1: Same as External 1, but the route was received through a not-so-stubby area (NSSA) • OSPF NSSA 2: Same as External 2, but the route was received through a not-so-stubby area <p>For more information about OSPF subtypes, see "OSPF" on page 2146.</p>
Network	The IP addresses and network masks of destination networks that the FortiGate can reach.
Gateway IP	The IP addresses of gateways to the destination networks.
Interfaces	The interface through which packets are forwarded to the gateway of the destination network.
Up Time	The total accumulated amount of time that a route learned through RIP, OSPF, or BGP has been reachable.
Distance	<p>The administrative distance associated with the route. A value of 0 means the route is preferable compared to other routes to the same destination, and the FortiGate may routinely use the route to communicate with neighboring routers and access servers.</p> <p>Modifying this distance for dynamic routes is route distribution. See "BGP" on page 2187.</p>

Field	Description
Metric	<p>The metric associated with the route type. The metric of a route influences how the FortiGate dynamically adds it to the routing table. The following are types of metrics and the protocols they are applied to:</p> <p>Hop count: Routes learned through RIP</p> <p>Relative cost: Routes learned through OSPF</p> <p>Multi-Exit Discriminator (MED): Routes learned through BGP. However, several attributes in addition to MED determine the best path to a destination network. For more information about BGP attributes, see "BGP" on page 2187. By default, the MED value associated with a BGP route is zero. However, the MED value can be modified dynamically. If the value was changed from the default, the Metric column displays a non-zero value.</p> <p>This field isn't displayed when IP version 6 is selected.</p>

Copying DSCP value in GRE tunnels

You can enable an option to allow copying of the DSCP (Differentiated services code point) value in GRE tunnels. This feature enables the keeping of the DSCP marking in the packets after encapsulation for going through GRE tunnels.

To enable DSCP copying - CLI:

```
config system gre-tunnel
  edit <name>
    set dscp-copying enable
  next
end
```

Configuring the maximum number of IP route cache entries

To configure the maximum number of route cache entries - CLI:

```
config system global
  set max-route-cache-size <number of cache entries>
end
```

where <number of cache entries> is in the range 0 to 2147483647

Unsetting the field causes the value to be set to the kernel-calculated default:

```
config system global
  unset max-route-cache-size
end
```

Viewing the routing table in the CLI

You can easily view the static routing table in the CLI. You can view the static routing table, just as in the GUI, or you can view the full routing table.

When you view the list of static routes using the `get router static` CLI command, the configured static routes are displayed. When you view the routing table using the `get router info routing-table all` CLI command, it's the entire routing table information that's displayed, including configured and learned routes of all types. The two commands show different information in different formats.



If VDOMs are enabled on the FortiGate, all routing-related CLI commands must be performed within a VDOM and not in the global context.

To view the routing table - CLI:

```
# get router info routing-table all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
* - candidate default

S* 0.0.0.0/0 [10/0] via 192.168.183.254, port2
S 1.0.0.0/8 [10/0] via 192.168.183.254, port2
S 2.0.0.0/8 [10/0] via 192.168.183.254, port2
C 10.142.0.0/23 is directly connected, port3
B 10.160.0.0/23 [20/0] via 10.142.0.74, port3, 2d18h02m
C 192.168.182.0/23 is directly connected, port2
```

Examining an entry:

```
B 10.160.0.0/23 [20/0] via 10.142.0.74, port3, 2d18h02m
```

Value	Description
B	BGP. The routing protocol used.
10.160.0.0/23	The destination of this route, including netmask.
[20/0]	20 indicates an administrative distance of 20 out of a range of 0 to 255. 0 is an additional metric associated with this route, such as in OSPF.
10.142.0.74	The gateway or next hop.
port3	The interface that the route uses.
2d18h02m	The age of the route (in this example, it's almost three days old).

To view the kernel routing table - CLI:

```
# get router info kernel

tab=254 vf=0 scope=253 type=1 proto=2 prio=0 0.0.0.0/0.0.0.0/0->10.11.201.0/24
pref=10.11.201.4 gwy=0.0.0.0 dev=5(external1)
```

```
tab=254 vf=0 scope=253 type=1 proto=2 prio=0 0.0.0.0/0.0.0.0/0->172.20.120.0/24
pref=172.20.120.146 gwy=0.0.0.0 dev=6(internal)
```

The parts of the routing table entry are:

Value	Description
tab	Table number: This will be either 254 (unicast) or 255 (multicast).
vf	Virtual domain of the firewall: This is the VDOM index number. If VDOMs aren't enabled, this number is 0.
type	Type of routing connection: Valid values include: 0 - unspecific 1 - unicast 2 - local 3 - broadcast 4 - anycast 5 - multicast 6 - blackhole 7 - unreachable 8 - prohibited
proto	Type of installation: This indicates where the route came from. Valid values include: 0 - unspecific 2 - kernel 11 - ZebOS routing module 14 - FortiOS 15 - HA 16 - authentication based 17 - HA1
prio	Priority of the route. Lower priorities are preferred.
->10.11.201.0/24 (->x.x.x.x/mask)	The IP address and subnet mask of the destination
pref	Preferred next hop along this route
gwy	Gateway: The address of the gateway this route will use
dev	Outgoing interface index: This number is associated with the interface for this route. If VDOMs are enabled, the VDOM is also included here. If an interface alias is set for this interface, it is also displayed here.

Searching the routing table

You can apply a filter to search the routing table and display only certain routes. For example, you can display one or more static routes, connected routes, routes learned through RIP, OSPF, or BGP, and routes associated with the network or gateway that you specify.

If you want to search the routing table by route type and further limit the display according to network or gateway, all of the values that you specify as search criteria must match corresponding values in the same routing table entry in order for that entry to be displayed. An implicit AND condition is applied to all of the search parameters you specify.

For example, if the FortiGate is connected to network 172.16.14.0/24 and you want to display all directly connected routes to network 172.16.14.0/24, you must select **Connected** from the **Type** list, type 172.16.14.0/24 in the **Network** field, and then select **Apply Filter** to display the associated routing table entry or entries. Any entry that contains the word "Connected" in its **Type** field and the specified value in the **Gateway** field will be displayed.

In this example, you will apply a filter to search for an entry for static route to 10.10.10.10/24.

To search the routing table routing table - GUI:

1. Go to **Monitor > Routing Monitor**.
2. From the **Type** list, select the type of route to display. In this example, select **Static**.
3. If you want to display routes to a specific network, type the IP address and netmask of the network in the **Networks** field. In our example, enter 10.10.10.10/24.
4. If you want to display routes to a specific gateway, type the IP address of the gateway in the **Gateway** field.
5. Select **Apply Filter**.



All of the values that you specify as search criteria must match corresponding values in the same routing table entry in order for that entry to be displayed.

To search the routing table - CLI:

```
FGT # get router info routing-table details 10.10.10.10
Routing entry for 10.10.10.10/24
  Known via "static", distance 10, metric 0, best
```

If there are multiple routes that match your filter, they will all be listed and the best match will be at the top of the list and indicated by the word "best".

Building the routing table

In factory default configuration, the routing table on the FortiGate contains a single static default route. You can add routing information to the routing table by defining additional static routes.

It's possible that the routing table is faced with several different routes to the same destination - the IP addresses of the next-hop router specified in those routes or the FortiGate interfaces associated with those routes may vary. In this situation, the best route is selected from the table.

The FortiGate selects the best route for a packet by evaluating the information in the routing table. The best route to a destination is typically associated with the shortest distance between the FortiGate and the closest

gateway, also known as a next-hop router. In some cases, the next best route may be selected if the best route is unavailable.

The FortiGate installs the best available routes in the forwarding table, which is a subset of the routing table. Packets are forwarded according to the information in the forwarding table.

Static routing security

Securing the information on your company network is a top priority for network administrators. Security is also required as the routing protocols used are internationally known standards that typically provide little or no inherent security by themselves.

The two reasons for securing your network are the sensitive and proprietary information on your network and also your external bandwidth. Hackers can steal not only your information, but they can also steal your bandwidth. Routing is a good low-level way to secure your network, even before UTM features are applied.

Routing provides security to your network in a number of ways including obscuring internal network addresses with NAT and blackhole routing, using RPF to validate traffic sources, and maintaining an access control list (ACL) to limit access to the network.

Network Address Translation

Network address translation (NAT) is a method of changing the address from which traffic appears to originate. This practice is used to hide the IP address on a company's internal networks and helps prevent malicious attacks that use those specific addresses.

This is accomplished by the router connected to that local network changing all the IP addresses to its externally connected IP address before sending the traffic out to the other networks, such as the Internet. Incoming traffic uses the established sessions to determine which traffic goes to which internal IP address. This also has the benefit of requiring only the router to be very secure against external attacks, instead of the entire internal network, as would be the case without NAT. Securing the network is much cheaper and easier to maintain.

Configuring NAT on a FortiGate includes the following steps:

1. Configure your internal network. For example, use the `10.11.101.0` subnet.
2. Connect your internal subnet to an interface on the FortiGate. For example, use `port1`.
3. Connect your external connection (for example, an ISP gateway of `172.20.120.2`) to another interface on the FortiGate (for example, `port2`).

Configure security policies to allow traffic between `port1` and `port2` on the FortiGate, ensuring that the NAT feature is enabled.

The above steps show that traffic from your internal network will originate on the `10.11.101.0` subnet and pass on to the `172.20.120.0` network. The FortiGate moves the traffic to the proper subnet. In doing that, the traffic appears to originate from the FortiGate interface on that subnet and it doesn't appear to originate from where it actually came from.

NAT "hides" the internal network from the external network. This provides security through obscurity. If a hacker tries to directly access your network, they will find the FortiGate, but they won't know about your internal network. The hacker would have to get past the security-hardened FortiGate to gain access to your internal network. NAT won't prevent hacking attempts that piggy back on valid connections between the internal network and the outside world. However, other UTM security measures can deal with these attempts.

Another security aspect of NAT is that many programs and services have problems with NAT. Consider if someone on the Internet tries to initiate a chat with someone on the internal network. The outsider can access

only the external interface on the FortiGate, unless the security policy allows the traffic through to the internal network. If allowed in, the correct internal user would respond to the chat. However, if it's not allowed, the request to chat will be refused or it will time out. This is accomplished in the security policy by allowing or denying different protocols.

Access control list

An access control list (ACL) is a table of addresses that have permission to send and receive data over a router's interface or interfaces. The router maintains an ACL, and when traffic comes in on a particular interface it's buffered, while the router checks the ACL to see if that traffic is allowed over that port. If it's allowed on that incoming interface, the next step is to check the ACL for the destination interface. If the traffic also passes that check, the buffered traffic is delivered to its destination. If either of those steps fail the ACL check, the traffic is dropped and an error message may be sent to the sender. The ACL ensures that traffic follows expected paths and any unexpected traffic isn't delivered. This stops many network attacks. However, to be effective, the ACL must be kept up to date. When employees or computers are removed from the internal network, their IP addresses must also be removed from the ACL. For more information about the ACL, see the router chapter of the [FortiOS CLI Reference](#).

Blackhole routes

A blackhole route is a route that drops all traffic sent to it. It is very much like `/dev/null` in Linux programming.

Blackhole routes are used to dispose of packets instead of responding to suspicious inquiries. This provides added security since the originator won't discover any information from the target network.

Blackhole routes can also limit traffic on a subnet. If some subnet addresses aren't in use, traffic to those addresses, which may be valid or malicious, can be directed to a blackhole for added security and to reduce traffic on the subnet.

The loopback interface, which is a virtual interface that doesn't forward traffic, was added to allow easier configuration of blackhole routing. Similar to a regular interface, the loopback interface has fewer parameters to configure and all traffic sent to it stops there. Since it can't have hardware connection or link status problems, it's always available, making it useful for other dynamic routing roles. Once configured, you can use a loopback interface in security policies, routing, and other places that refer to interfaces. You configure this feature only from the CLI. For more information, see the system chapter of the [FortiOS CLI Reference](#).

Configuring IPv6 blackhole routes

You can configure IPv6 blackhole routes. In the FortiGate GUI, select **Network > Static Routes** and select **Create New**. In the **Interface** field, choose **Blackhole**.

New Static Route

Destination IP/Mask

Interface

☐ Blackhole

Administrative Distance

Comments

0/255

Status

Enabled

Disabled

OK

Cancel

Blackhole static routing

System administrators use blackhole routing to divert unwanted traffic, such as packets from a Denial of Service (DoS) attack or communications from an illegal source. The traffic is routed to a dead interface, or a host designed to collect information for investigation. This mitigates the impact of the attack on the network.

To enable blackhole routing - CLI:

```
config router {static|static6}
  edit <sequence number>
    set blackhole enable
  next
end
```

Blackhole route priority

You can add a priority to a blackhole route to change its position relative to kernel routes in the routing table.

To add a blackhole route with a priority - CLI:

```
config router static
  edit <sequence number>
    set blackhole enable
    set priority 200
  next
end
```

Static routes and VRFs

You can configure static route support for multiple virtual routing and forwarding (VRFs) on a FortiGate.

To add VRFs for blackhole routes - CLI:

```
config router static
  edit <sequence-number>
    set vrf <VRF-ID>
  end
```

where `vrf` is a value of 0 to 31

Reverse path lookup

Whenever a packet arrives at one of the interfaces on the FortiGate, the FortiGate determines whether the packet was received on a legitimate interface by doing a reverse lookup using the source IP address in the packet header. This is also called anti-spoofing. If the FortiGate can't communicate with the computer at the source IP address through the interface on which the packet was received, the FortiGate drops the packet as it's likely a hacking attempt.

If the destination address can be matched to a local address, and the local configuration permits delivery, the FortiGate delivers the packet to the local network. If the packet is destined for another network, the Fortigate forwards the packet to a next-hop router according to a policy route and the information stored in the FortiGate forwarding table.

Multipath routing and determining the best route

Multipath routing occurs when more than one entry to the same destination is present in the routing table. When multipath routing happens, the FortiGate may have several possible destinations for an incoming packet, forcing the FortiGate to decide which next-hop is the best one.

It should be noted that some IP addresses will be rejected by routing protocols. These are called Martian addresses. They are typically IP addresses that are invalid and not routable because they have been assigned an address by a misconfigured system, or are spoofed addresses.

Two methods to manually resolve multiple routes to the same destination are to lower the administrative distance of one route or to set the priority of both routes. For the FortiGate to select a primary (preferred) route, manually lower the administrative distance associated with one of the possible routes. Setting the priority on the routes is a FortiGate feature and may not be supported by routers that aren't Fortinet products.

Administrative distance is based on the expected reliability of a given route. It's determined through a combination of the number of hops from the source and the protocol used. A hop is when traffic moves from one router to the next. More hops from the source means more possible points of failure. The administrative distance can be in the range of 1 to 255, with lower numbers being preferred. A distance of 255 is seen as infinite and won't be installed in the routing table.

Here's an example to illustrate how administration distance works. If there are two possible routes traffic can take between two destinations, with administration distances of 5 (always up) and 31 (sometimes not available), the traffic will use the route with an administrative distance of 5. If for some reason the preferred route (admin distance of 5) isn't available, the other route will be used as a backup.

Different routing protocols have different default administrative distances. These different administrative distances are based on a number of factors of each protocol such as reliability, speed, and so on. You can configure the default administrative distances for any of these routing protocols.

Default administrative distances for routing protocols and connections

Routing protocol	Default administrative distance
Direct physical connection	1
Static	10
EBGP	20
OSPF	110
IS-IS	115
RIP	120
IBGP	200

Another method to determine the best route is to manually change the priority of both routes in question. If the next-hop administrative distances of two routes on the FortiGate are equal, it may not be clear which route the packet will take. Manually configuring the priority for each of those routes will make it clear which next-hop will be used in the case of a tie. The priority for a route can be set in the CLI, or when editing a specific static route, as

described in the next section. Lower priority routes are preferred. Priority is a Fortinet value that may or may not be present in other brands of routers.

All entries in the routing table are associated with an administrative distance. If the routing table contains several entries that point to the same destination (the entries may have different gateways or interface associations), the FortiGate compares the administrative distances of those entries first, selects the entries having the lowest distances, and installs them as routes in the FortiGate forwarding table. As a result, the FortiGate forwarding table contains only those routes that have the lowest distances to every possible destination. While only static routing uses administrative distance as its routing metric, other routing protocols, such as RIP, can use metrics that are similar to administrative distance.

Route priority

After a FortiGate selects static routes for the forwarding table based on their administrative distances, the priority field of those routes determines routing preference. Priority is a Fortinet value that may or may not be present in other brands of routers.

You can configure the priority field through the CLI or the GUI. Priority values can range from 0 to 4 294 967 295. The route with the lowest value in the priority field is considered the best route. It's also the primary route.

To change the priority of a route - GUI:

1. Go to **Network > Static Routes**.
2. Select the route entry, and select **Edit**.
3. Select **Advanced Options**.
4. Enter the **Priority** value.
5. Select **OK**.

To change the priority of a route - CLI:

The following command changes the priority to 5 for a route to the address 10.10.10.1 on the port1 interface.

```
config router static
  edit 1
    set device port1
    set gateway 10.10.10.10
    set dst 10.10.10.1
    set priority 5
  next
end
```

If there are other routes set to priority 10, the route set to priority 5 will be preferred. If there are routes set to priorities less than 5, those other routes will be preferred instead.

In summary, because you can use the CLI to specify which sequence numbers or priority field settings to use when defining static routes, you can prioritize routes to the same destination according to their priority field settings. For a static route to be the preferred route, you must create the route using the `config router static` CLI command and specify a low priority for the route. If two routes have the same administrative distance and the same priority, then they are equal-cost multi-path (ECMP) routes.

Since this means there is more than one route to the same destination, it can be confusing which route or routes to install and use. However, if you have enabled load balancing with ECMP routes, different sessions will resolve this problem by using different routes to the same address.

Use of firewall addresses for static route destinations

To help prevent false positive when scanning for duplicate static routes, the `dst_addr` field is also checked.

Removing RPF checks from the state evaluation process

You can remove RPF (reverse path forwarding) state checks without needing to enable asymmetric routing. You can disable state checks for traffic received on specific interfaces.



Disabling state checks makes a FortiGate less secure and should only be done with caution.

To remove RPF checks from the state evaluation process - CLI:

```
config system interface
  edit <interface_name>
    set src-check disable
  next
end
```

Troubleshooting static routing

When there are problems with your network that you think are related to static routing, there are a few basic tools available to locate the problem. These tools include ping, traceroute, and examining routing table contents.

Ping

Beyond the basic connectivity information, ping can tell you the amount of packet loss (if any), how long it takes the packet to make the round trip, and the variation in that time from packet to packet.

If there's no packet loss detected, your basic network connectivity is okay.

If there's some packet loss detected, you should investigate:

- Possible ECMP, split horizon, network loops
- Cabling to ensure no loose connections

If there's total packet loss, you should investigate:

- Hardware: Ensure cabling is correct, and all equipment between the two locations is accounted for
- Addresses and routes: Ensure all IP addresses and routing information along the route is configured as expected
- Firewalls: Ensure all firewalls are set to allow PING to pass through

To ping from a Windows PC:

1. Go to a DOS prompt. Typically you go to **Start > Run**, enter `cmd` and select **OK**.
2. Enter `ping 10.11.101.100` to ping the default internal interface of the FortiGate with four packets.

To ping from an Apple computer:

1. Open the Terminal.
2. Enter `ping 10.11.101.100`.
3. If the ping fails, it will stop after a set number of attempts. If it succeeds, it will continue to ping repeatedly. Press `Control+C` to end the attempt and see gathered data.

To ping from a Linux PC:

1. Go to a command line prompt.
2. Enter `"/bin/etc/ping 10.11.101.101"`.

Traceroute

Where ping will only tell you if it reached its destination and came back successfully, traceroute will show each step of its journey to its destination and how long each step takes. If ping finds an outage between two points, you can use traceroute to locate exactly where the problem is.

To use traceroute on a Windows PC:

1. Go to a DOS prompt. Typically you go to **Start > Run**, enter `cmd` and select **OK**.
2. Enter `"tracert fortinet.com"` to trace the route from the PC to the Fortinet website.

To use traceroute from an Apple computer:

1. Open the Terminal.
2. Enter `traceroute fortinet.com`.
3. The terminal will list the number of steps made. Upon reaching the destination, it will list three asterisks per line. Press `Control+C` to end the attempt.

To use traceroute on a Linux PC:

1. Go to a command line prompt.
2. Enter `"/bin/etc/traceroute fortinet.com"`.
The Linux traceroute output is very similar to the MS Windows traceroute output.

Examine routing table contents

The first place to look for information is the routing table.

The routing table is where all the currently used routes are stored for both static and dynamic protocols. If a route is in the routing table, it saves the time and resources of a lookup. If a route isn't used for a while and a new route needs to be added, the oldest least used route is bumped if the routing table is full. This ensures the most recently used routes stay in the table. Note that if a FortiGate is in transparent mode, you won't be able to perform this step.

If the FortiGate is running in NAT mode, verify that all desired routes are in the routing table: local subnets, default routes, specific static routes, and dynamic routing protocols.

To check the routing table in the GUI, use the Routing Monitor. Go to **Monitor > Routing Monitor**. In the CLI, use the `get router info routing-table all` command.

Static routing tips

When your network goes beyond basic static routing, here are some tips to help you plan and manage your static routing.

Always configure a default route

The first thing you configure on a router on your network should be the default route. And where possible, the default routes should point to either one or very few gateways. This makes it easier to locate and correct problems in the network. By comparison, if one router uses a second router as its gateway which uses a fourth for its gateway and so on, one failure in that chain will appear as an outage for all the devices downstream. By using one or very few addresses as gateways, if there's an outage on the network, it will either be very localized or network-wide. Either outage is easy to troubleshoot.

Have an updated network plan

A network plan lists different subnets, user groups, and different servers. Essentially, it puts all your resources on the network and shows how the parts of your network are connected. Keeping your plan updated will also help you troubleshoot problems more quickly when they arise.

A network plan helps your static routing by eliminating potential bottlenecks and helping troubleshoot any routing problems that come up. Also, you can use it to plan for the future and act on any changes to your needs or resources more quickly.

Plan for expansion

No network remains the same size. At some time, all networks grow. If you take future growth into account, there will be less disruption to your existing network when that growth happens. For example, allocating a block of addresses for servers can easily prevent having to re-assign IP addresses to multiple servers due to a new server.

With static routing, if you group parts of your network properly you can easily use network masks to address each part of your network separately. This will reduce the amount of administration required both to maintain the routing and to troubleshoot any problems.

Configure as much security as possible

Securing your network through static routing methods is a good low level method to defend both your important information and your network bandwidth.

- Implement NAT to obscure your IP address is an excellent first step
- Implement blackhole routing to hide which IP addresses are in use or not on your local network
- Configure and use access control list (ACL) to help ensure you know only valid users are using the network

All three features limit access to the people who should be using your network and obscure your network information from the outside world and potential hackers.

Policy routing

Policy routing allows you to redirect traffic away from a static route. This can be useful if you want to route certain types of network traffic differently. You can use incoming traffic's protocol, source address or interface, destination address, or port number to determine where to send the traffic. For example, generally network traffic

will go to the router of a subnet, but you might want to direct SMTP or POP3 traffic directly to the mail server on that subnet.

If you configure the FortiGate with routing policies and a packet arrives at the FortiGate, the FortiGate starts at the top of the Policy Route list and attempts to match the packet with a policy. If a match is found and the policy contains enough information to route the packet (a minimum of the IP address of the next-hop router and the FortiGate interface for forwarding packets to it), the FortiGate routes the packet using the information in the policy. If no policy route matches the packet, the FortiGate routes the packet using the routing table.



Most policy settings are optional, and a matching policy alone might not provide enough information for forwarding the packet. In fact, the FortiGate almost always requires a matching route in the routing table in order to use a policy route. The FortiGate refers to the routing table in an attempt to match the information in the packet header with a route in the routing table.

Policy route options define which attributes of an incoming packet cause policy routing to occur. If the attributes of a packet match all the specified conditions, the FortiGate routes the packet through the specified interface to the specified gateway.

To view policy routes go to **Network > Policy Routes**.

Field	Description
Create New	Add a policy route. See "Adding a policy route" on page 2079 .
Edit	Edit the selected policy route.
Delete	Delete the selected policy route.
Move To	Move the selected policy route. Enter the new position and select OK . For more information, see "Moving a policy route" on page 2082 .
#	The ID numbers of configured route policies. These numbers are sequential unless policies have been moved within the table.
Incoming	The interfaces on which packets subjected to route policies are received.
Outgoing	The interfaces through which policy routed packets are routed.
Source	The IP source addresses and network masks that cause policy routing to occur.
Destination	The IP destination addresses and network masks that cause policy routing to occur.

Adding a policy route

To add a policy route, go to **Network > Policy Routes** and select **Create New**.

Field	Description
Protocol	<p>Select from existing or specify the protocol number to match. The Internet Protocol Number is found in the IP packet header. RFC 5237 describes protocol numbers and you can find a list of the assigned protocol numbers here. The range is from 0 to 255. A value of 0 disables the feature.</p> <p>Commonly used Protocol settings include 6 for TCP sessions, 17 for UDP sessions, 1 for ICMP sessions, 47 for GRE sessions, and 92 for multicast sessions.</p>
Incoming Interface	Select the name of the interface through which incoming packets subjected to the policy are received.
Source Address / Mask	To perform policy routing based on IP source address, type the source address and network mask to match. A value of 0.0.0.0/0.0.0.0 disables the feature.
Destination Address / Mask	To perform policy routing based on the IP destination address of the packet, type the destination address and network mask to match. A value of 0.0.0.0/0.0.0.0 disables the feature.
Destination Ports	<p>To perform policy routing based on the port on which the packet is received, type the same port number in the From and To fields. To apply policy routing to a range of ports, type the starting port number in the From field and the ending port number in the To field. A value of 0 disables this feature.</p> <p>The Destination Ports fields are only used for TCP and UDP protocols. The ports are skipped over for all other protocols.</p>
Type of Service	Use a two digit hexadecimal bit pattern to match the service, or use a two digit hexadecimal bit mask to mask out. For more information, see "Type of service" on page 2081 .
Outgoing Interface	Select the name of the interface through which packets affected by the policy will be routed.
Gateway Address	Type the IP address of the next-hop router that the FortiGate can access through the specified interface.

Example policy route

Configure the following policy route to send all FTP traffic received at `port1` out the `port10` interface and to a next hop router at IP address `172.20.120.23`. To route FTP traffic, set protocol to 6 (for TCP) and set both of the destination ports to 21 (the FTP port).

Field	Value
Protocol	6

Field	Value
Incoming interface	port1
Source address / mask	0.0.0.0/0.0.0.0
Destination address / mask	0.0.0.0/0.0.0.0
Destination Ports	From 21 to 21
Type of Service	bit pattern: 00 (hex) bit mask: 00 (hex)
Outgoing interface	port10
Gateway Address	172.20.120.23

Enabling or disabling individual policy routes

You can enable or disable individual policy routes.

To configure IPv4 policy routes - CLI:

```
config router policy
  edit <sequence number>
    set status {enable | disable}
  next
end
```

To configure IPv6 policy routes - CLI:

```
config router policy6
  edit <sequence number>
    set status {enable | disable}
  next
end
```

Type of service

Type of service (TOS) is an 8-bit field in the IP header that allows you to determine how the IP datagram should be delivered, with qualities, such as delay, priority, reliability, and minimum cost.

Each quality helps gateways determine the best way to route datagrams. A router maintains a ToS value for each route in its routing table. The lowest priority TOS is 0, the highest is 7 - when bits 3, 4, and 5 are all set to 1. The router tries to match the TOS of the datagram to the TOS on one of the possible routes to the destination. If there's no match, the datagram is sent over a zero TOS route.

Using increased quality may increase the cost of delivery because better performance may consume limited network resources. For more information, see [RFC 791](#) and [RFC 1349](#).

The role of each bit in the IP header TOS 8-bit field

Bit	Quality	Description
bits 0, 1, 2	Precedence	Some networks treat high precedence traffic as more important traffic. Precedence should only be used within a network, and can be used differently in each network. Typically, you don't care about these bits.
bit 3	Delay	When set to 1, this bit indicates low delay is a priority. This is useful for such services as VoIP where delays degrade the quality of the sound.
bit 4	Throughput	When set to 1, this bit indicates high throughput is a priority. This is useful for services that require lots of bandwidth, such as video conferencing.
bit 5	Reliability	When set to 1, this bit indicates high reliability is a priority. This is useful when a service must always be available, such as with DNS servers.
bit 6	Cost	When set to 1, this bit indicates low cost is a priority. Generally there is a higher delivery cost associated with enabling bits 3, 4, or 5, and bit 6 indicates to use the lowest cost route.
bit 7	Reserved for future use	Not used at this time.

For example, if you want to assign low delay and high reliability for a VoIP application, where delays are unacceptable, you would use a bit pattern of xxx1x1xx where 'x' indicates that bit can be any value. Since all bits aren't set, this is a good use for the bit mask. If the mask is set to 0x14, it will match any TOS packets that are set to low delay and high reliability.

Moving a policy route

A routing policy is added to the bottom of the routing table when it's created. If you prefer to use one policy over another, you may want to move it to a different location in the routing policy table.

The option to use one of two routes happens when both routes are a match, for example 172.20.0.0/255.255.0.0 and 172.20.120.0/255.255.255.0. If both of these routes are in the policy table, both can match a route to 172.20.120.112 but you would consider the second one a better match. In this case, the best match route should be positioned before the other route in the policy table.

To change the position of a policy route in the table, go to **Network > Policy Routes** and select **Move To** for the policy route you want to move.

Field	Description
Before/After	Select Before to place the selected policy route before the indicated route. Select After to place it following the indicated route.
Policy route ID	Enter the policy route ID of the route in the policy route table to move the selected route before or after.

Use of firewall addresses for policy route destinations

When you configure a policy route, you can use firewall addresses and address groups. The only exception for the address types that can be used is the URL type of address object.

Static routing in transparent mode

FortiOS operating modes allow you to change the configuration of a FortiGate, depending on its role in your network. A FortiGate can operate in two different modes: transparent operation mode or NAT operation mode. In transparent mode, a FortiGate acts as a bridge where all physical interfaces act like one interface. A FortiGate broadcasts traffic that arrives through any interface out through all interfaces.

In transparent mode, you install a FortiGate between your internal network and your router. The FortiGate doesn't make any changes to IP addresses and applies only security scanning to traffic. When you add a FortiGate to a network in transparent mode, you don't have to make any network changes except configuring the FortiGate with a management IP address. You usually use transparent mode when you want to increase your network protection but it's impractical to change your network configuration.

When you configure routing in transparent mode on a FortiGate, all interfaces must be connected to the same subnet. This means all traffic comes from and leaves on the same subnet. This is important because it limits the static routing options to only gateways that are attached to this subnet. For example, if you have only one router that connects your network to the Internet, all static routing on the FortiGate uses this gateway. For this reason, static routing on a FortiGate in transparent mode may be a bit different, but it's not as complex as routing in NAT mode.

To view the routing table in transparent mode, go to **Network > Static Routes**. When you view entries for static routes in transparent mode, you'll see the following settings:

Field	Description
Destination	When Subnet is selected, shows the IP address and netmask of the destination of the traffic being routed. 0.0.0.0 is the default route and matches all traffic destinations.
Gateway	Specifies the IP address of the next hop for traffic. This is usually the IP address of a router on the edge of your network.
Priority	<p>The FortiGate uses the priority if there's more than one match for a route. This allows you to use multiple routes, but configure preferred routes.</p> <p>Routes with a larger value have a lower priority. If the preferred route isn't available, another route is used instead. If there is more than one match for a route, and the routes have the same priority, the FortiGate uses Equal Cost Multiple Path (ECMP) to share traffic between the routes.</p> <p>The possible values are 0 to 4294967295. This setting only applies to static routes. The priority for routes that are dynamically learned from routing protocols is 0.</p>

For more information about configuring a FortiGate in transparent mode, see the [FortiOS Transparent Mode Handbook](#).

Source prefixes for static routes in transparent mode

If a FortiGate has more than one management IP address and default route, packets can't differentiate between them and may reach the wrong management IP address. To avoid this, you can configure a source prefix that allows the FortiGate to differentiate between multiple default routes. This is necessary only for static routes in transparent mode.

To configure source prefixes - CLI:

```
config router static
  edit <sequence-number>
    set gateway <IP-address>
    set src <source-prefix>
  next
  edit <sequence-number>
    set gateway <IP-address>
    set src <source-prefix>
  next
end
```



This command is only available in transparent mode.

Static routing example

This is an example of a typical small network configuration that uses only static routing.

This network is in a dental office that includes a number of dentists, assistants, and office staff. The size of the office isn't expected to grow significantly in the near future, and the network usage is very stable (there are no new applications being added to the network).

The users on the network are:

- Administrative staff: Access to local patient records to perform online billing
- Dentists: Access and update local patient records to research online from desk
- Assistants: Access and update local patient records in exam rooms

The distinction here is mainly that only the administrative staff and dental office need access to the Internet. All other traffic is local and doesn't need to leave the local network. Routing is only required for the outbound traffic and the computers that have valid outbound traffic.



Configuring routing only on computers that need it acts as an additional layer of security by helping prevent malicious traffic from leaving the network.

Network layout and assumptions

The computers on the network are administrative staff computers, dental office computers, and dental exam room computers. While there are other devices on the local network, such as printers, they don't need Internet access or any routing.

The networked office equipment includes one PC for administrative staff, 3 PCs for dentists, and 5 PCs in the exam rooms. There's also a network printer and a router on the network.

Assumptions about these computers and network include:

- The FortiGate is a model with interfaces labeled port1 and port2.
- The FortiGate has been installed and is configured in NAT/Route mode.
- VDOMs aren't enabled.
- The computers on the network are running MS Windows software.
- Any hubs required in the network aren't shown in the network diagram.
- The network administrator has access to the ISP IP addresses and is the super_admin administrator on the FortiGate.

Static routing example device names, IP addresses, and level of access

Device name	IP address	Need external access?
Router	192.168.10.1	Yes
Admin	192.168.10.11	Yes
Dentist1-3	192.168.10.21-23	Yes
Exam1-5	192.168.10.31-35	No
Printer	192.168.10.41	No

General configuration steps

The steps to configuring routing on this network are:

1. [Get your ISP information such as DNS, gateway, etc.](#)
2. [Configure the FortiGate](#)
3. [Configure the PCs for the administrator and dentists](#)
4. [Testing network configuration](#)

Get your ISP information such as DNS, gateway, etc.

Your local network connects to the Internet through your Internet Service Provider (ISP). They have IP addresses that you need to configure your network and routing.

The addresses that you need for routing are your assigned IP address, DNS servers, and the gateway.

Configure the FortiGate

The FortiGate has two interfaces in use: one connected to the internal network and one connected to the external network. Port1 is the internal interface and port2 is the external interface.

To configure the FortiGate:

1. [Configure the internal interface \(port1\)](#)
2. [Configure the external interface \(port2\)](#)
3. [Configure networking information](#)

4. [Configure basic security policies](#)
5. [Configure static routing](#)

Configure the internal interface (port1)

To configure the internal interface (port1) - GUI:

1. Go to **Network > Interfaces**. Highlight **port1** and select **Edit**.
2. Enter the following information:

Addressing Mode	Manual
IP/Netmask	172.100.1.1/255.255.255.0
Administrative Access	HTTPS, PING, TELNET
Description	Internal network

To configure the internal interface (port1) - CLI:

```
config system interface
  edit port1
    set IP 192.168.10.1 255.255.255.0
    set allowaccess https ping telnet
    set description "internal network"
  next
end
```

Configure the external interface (port2)

The external interface connects to your ISP network. You need to know the IP addresses in their network that you should connect to. In this example, the address that the ISP gave you is 172.100.20.20, which will connect to the gateway at 172.100.20.5 on their network, and their DNS servers are 172.11.22.33 and 172.11.22.34.

To configure the internal interface (port2) - GUI:

1. Go to **Network > Interfaces**. Highlight **port2** and select **Edit**.
2. Enter the following:

Addressing Mode	Manual
IP/Netmask	172.100.20.20/255.255.255.0
Administrative Access	HTTPS, PING, TELNET
Description	Internal network

To configure the internal interface (port2) - CLI:

```
configure system interface
  edit port2
    set IP 172.100.20.20 255.255.255.0
    set allowaccess https ping telnet
```

```
        set description "internal network"
    next
end
```

Configure networking information

Networking information includes the gateway and DNS servers. A FortiGate requires a connection to the Internet for antivirus and other periodic updates.

To configure networking information - GUI:

1. Go to **Network > DNS**.
2. Enter the primary and secondary DNS addresses.
3. Select **Apply**.

To configure networking information - CLI:

```
config system global
    set dns_1 172.11.22.33
    set dns_2 172.11.22.34
end
```

Configure basic security policies

For traffic to flow between the internal and external ports in both directions, as a minimum, two security policies are required. More can be used to further limit or direct traffic, as needed, but won't be included here.

Before configuring the security policies, a firewall address group is configured for the PCs that are allowed Internet access. This prevents a PC without Internet privileges from accessing the Internet.

The security policy assumptions are:

- For added security, only the basic networking services are listed as allowed. Others can easily be added as users require them.
- In this example, to keep things simple, both incoming and outgoing security policies are the same. In a real network there are applications that are allowed out but not in, and vice versa.
- Endpoint control has been enabled to ensure that all computers on the local network are running FortiClient and those installs are up to date. This feature ensures added security on your local network without the need for the network administrator to continually bother users to update their software. The FortiGate can store an up to date copy of the FortiClient software and offer a URL to it for users to install it if they need to.

To configure security policies - GUI:

1. Go to **Policy & Objects > Objects > Addresses**.
2. Create a new Firewall Address entry for each of:

PC Name	IP Address	Interface
Admin	192.168.10.11	port1
Dentist1	192.168.10.21	port1

Dentist2	192.168.10.22	port1
Dentist3	192.168.10.23	port1

3. Go to **Policy & Objects > Objects > Addresses**.
4. Select the dropdown arrow next to **Create New** and select **Address Group**.
5. Name the group Internet_PCs.
6. Add Admin, Dentist1, Dentist2, and Dentist3 as members of the group.
7. Select **OK**.
8. Go to **Policy & Objects > Policy > IPv4**.
9. Select **Create New**.
10. Enter the following: DH - port2(external) -> port1(internal)

Incoming Interface	port2
Source Address	all
Outgoing Interface	port1
Destination Address	Internet_PCs
Schedule	always
Service	Multiple. Select DHCP, DNS,FTP, HTTP, HTTPS, NTP, POP3, SMTP, SSH.
Action	ACCEPT
Log Allowed Traffic	Enabled

11. Select **OK**.
12. Select **Create New**.
13. Enter the following:

Incoming Interface	port1
Source Address	Internet_PCs
Outgoing Interface	port2
Destination Address	all
Schedule	always

Service	Multiple. Select DHCP, DNS,FTP, HTTP, HTTPS, NTP, POP3, SMTP, SSH.
Action	ACCEPT
Log Allowed Traffic	Enabled

14. Select **OK**.

To configure security policies - CLI:

```
config firewall address
  edit "Admin"
    set associated-interface "port1"
    set subnet 192.168.10.11 255.255.255.255
  next
  edit "Dentist1"
    set associated-interface "port1"
    set subnet 192.168.10.21 255.255.255.255
  next
  edit "Dentist2"
    set associated-interface "port1"
    set subnet 192.168.10.22 255.255.255.255
  next
  edit "Dentist3"
    set associated-interface "port1"
    set subnet 192.168.10.23 255.255.255.255
  end
config firewall addrgrp
  edit Internet_PCs
    set member Admin Dentist1 Dentist2 Dentist3
  end
config firewall policy
  edit 1
    set srcintf port1
    set dstintf port2
    set srcaddr Internet_PCs
    set dstaddr all
    set action accept
    set schedule always
    set service "DHCP" "DNS" "FTP" "HTTP" "HTTPS" "NTP" "POP3" "SMTP" "SSH"
    set logtraffic enable
    set label "Section2"
    set endpoint-restrict-check no-av db-outdated
  next
  edit 2
    set srcintf port2
    set dstintf port1
    set srcaddr all
    set dstaddr Internet_PCs
    set action accept
    set schedule always
    set service "DHCP" "DNS" "FTP" "HTTP" "HTTPS" "NTP" "POP3" "SMTP" "SSH"
    set logtraffic enable
```



```

        set label "Section2"
        set endpoint-restrict-check no-av db-outdated
    next
end

```

Adding FortiClient enforcement to interfaces

You can enforce the use of FortiClient on individual interfaces.

In the FortiGate GUI, select **Network > Interfaces** and choose an interface. Under the **Admission Control** heading, you can enable the **Allow FortiClient Connections** setting. Once you enable this setting, two more options become visible: **Discover Clients (Broadcast)** and **FortiClient Enforcement**. When you enable FortiClient enforcement, you enforce that in order for incoming traffic to pass through that interface, it must be initiated by a device running FortiClient.

Once you enforce the use of FortiClient on the interface, you should also configure FortiClient profiles for the incoming connections. You can also set up any exemptions that are needed. Just below the **FortiClient Enforcement** option are fields for **Exempt Sources** and **Exempt Destinations/Services**. These can be selected from address or services objects already configured on the FortiGate.

In the CLI, use the following commands:

```

config system interface
    edit port1
        set listen-forticlient-connection [enable|disable]
        set endpoint-compliance [enable|disable]
    next
end

```

Configure static routing

With the rest of the FortiGate configured, static routing is the last step before moving on to the rest of the local network. All traffic on the local network will be routed according to this static routing entry.

To configure Fortinet static routing - GUI:

1. Go to **Network > Static Routes**.
2. Select the top route on the page and then select **Edit**.
3. Enter the following information:

Destination IP/Mask	172.100.20.5
Gateway	172.100.20.5
Interface	port2
Administrative Distance	10

4. Select **OK**.

To configure Fortinet unit static routing - CLI:

```

configure routing static
    edit 1

```

```
set gateway 172.100.20.5
set distance 10
set device port2
set dst 0.0.0.0
next
end
```

Configure the PCs for the administrator and dentists

After the router is configured, we need to configure the computers that require Internet access. These computers need routing to be configured on them. As the other computers don't require routing, they aren't included here.

The procedure to configure these computers is the same. Repeat the following procedure for the corresponding PCs.



The Windows CLI procedure doesn't configure the DNS entries. It just adds the static routes.

To configure routing and DNS on PCs for administrator and dentists - Windows GUI:

1. On the PC, select **Start > Control Panel > Network Connections**.
2. Right click on the network connection to your local network that has a status of Connected, and select **Properties**.
3. Under the **General** tab, from the list select **TCP/IP**, and **Properties**.
4. Under **Gateway**, enter the FortiGate unit address (192.168.10.1).
5. Enter the primary and secondary DNS server addresses from your ISP (172.11.22.33 and 172.11.22.34).
6. Select **OK**.

To configure routing on PCs for administrator and dentists - Windows CLI:

1. On the PC, select **Start > Run**, enter "cmd", and select **OK**.
2. At the command prompt, type:

```
route ADD 0.0.0.0 MASK 0.0.0.0 172.100.20.5 METRIC 10
route ADD 192.168.10.0 MASK 255.255.255.0 192.168.10.1 METRIC 5
```
3. Confirm these routes have been added. Type:

```
route PRINT
```

If you don't see the two routes you added, try adding them again, while paying attention to avoid spelling mistakes.
4. Test that you can communicate with other computers on the local network, and with the Internet. If there are no other computers on the local network, connect to the FortiGate.

Configure other PCs on the local network

The PCs on the local network without Internet access (for example, the exam room PCs) can be configured now.

As this step doesn't require any routing, details haven't been included.

Testing network configuration

There are three tests that you can run on the network to ensure proper connectivity:

- Test that PCs on the local network can communicate
- Test that Internet_PCs on the local network can access the Internet
- Test that non-Internet_PCs can't access the Internet

Test that PCs on the local network can communicate

1. Select any two PCs on the local network, such as Exam4 and Dentist3.
2. On the Exam4 PC, at the command prompt, enter `ping 192.168.10.23`.

The output from this command should appear similar to the following:

```
Pinging 192.168.10.23 with 32 bytes of data:
```

```
Reply from 192.168.10.23: bytes=32 time<1m TTL=255
Reply from 192.168.10.23: bytes=32 time<1m TTL=255
Reply from 192.168.10.23: bytes=32 time<1m TTL=255
```

3. At the command prompt, enter `exit` to close the window.
4. On the Dentist3 PC, at the command prompt, enter `ping 192.168.10.34`.

The output from this command should appear similar to the following:

```
Pinging 192.168.10.34 with 32 bytes of data:
```

```
Reply from 192.168.10.34: bytes=32 time<1m TTL=255
Reply from 192.168.10.34: bytes=32 time<1m TTL=255
Reply from 192.168.10.34: bytes=32 time<1m TTL=255
```

5. At the command prompt, enter `exit` to close the window.
6. Repeat these steps for all PCs on the local network.

If the output doesn't appear similar to above, there's a problem with the network configuration between these two PCs.

To test that Internet_PCs on the local network can access the Internet

The easiest way to access the Internet is with an Internet browser. However, if that doesn't work, it's best to do a traceroute to see at what point the problem is. This can help determine if it's a networking problem, such as cabling, or if it's an access problem, such as this PC not having Internet access.

1. Select any PC on the local network that's supposed to have Internet access, such as Admin.
2. On the Admin PC, open an Internet browser and attempt to access a website on the Internet, such as <http://www.fortinet.com>.

If this is successful, this PC has Internet access.

3. If step2 wasn't successful, at the command prompt on the PC, enter `tracert 22.11.22.33`.

The output from this command should appear similar to:

```
Pinging 22.11.22.33 with 32 bytes of data:
```

```
Reply from 22.11.22.33: bytes=32 time<1m TTL=255
Reply from 22.11.22.33: bytes=32 time<1m TTL=255
Reply from 22.11.22.33: bytes=32 time<1m TTL=255
```

Dynamic routing

Dynamic routing uses a dynamic routing protocol to automatically select the best route to put into the routing table. Instead of having to manually enter static routes in the routing table, dynamic routing automatically receives routing updates and dynamically decides which routes are best to go into the routing table. It's this intelligent and hands-off approach that makes dynamic routing so useful.

Dynamic routing protocols vary in many ways and this is reflected in the various administrative distances assigned to routes learned from dynamic routing. These variations take into account differences in reliability, speed of convergence, and other similar factors. For more information about these administrative distances, see ["Advanced static routing" on page 2062](#).

Overview

Comparing static and dynamic routing

A common term used to describe dynamic routing is convergence. Convergence is the ability to work around network problems and outages, for the routing to come together despite obstacles. For example, if the main router between two endpoints goes down, convergence is the ability to find a way around that failed router and reach the destination. Static routing has zero convergence beyond trying the next route in its limited local routing table. If a network administrator doesn't fix a routing problem manually, it may never be fixed and may result in a downed network. Dynamic routing solves this problem by involving routers along the route in the decision-making process about the optimal route, and using the routing tables of these routers to find potential routes around the outage. In general, dynamic routing has better scalability, robustness, and convergence. However, the cost of these added benefits includes more complexity and some overhead. For example, the routing protocol uses some bandwidth for its own administration.

Comparing static and dynamic routing

Feature	Static routing	Dynamic routing
Hardware support	Supported by all routing hardware	May require special, more expensive routers
Router memory required	Minimal	Can require considerable memory for larger tables
Complexity	Simple	Complex
Overhead	None	Varying amounts of bandwidth used for routing protocol updates
Scalability	Limited to small networks	Very scalable, better for larger networks
Robustness	None: if a route fails, it has to be fixed manually	Robust: traffic routed around failures automatically

Feature	Static routing	Dynamic routing
Convergence	None	Varies from good to excellent

Dynamic routing protocols

A dynamic routing protocol is an agreed-on method of routing that the sender, receiver, and all routers along the path (route), support. Typically, the routing protocol involves a process running on all computers and routers along that route to enable each router to handle routes in the same way as the others. The routing protocol determines how the routing tables are populated along that route, how the data is formatted for transmission, and what information about a route is included with that route. For example, RIP and BGP use distance vector algorithms and OSPF uses a shortest path first algorithm. Each routing protocol has different strengths and weaknesses. One protocol may have fast convergence, while another may be very reliable, and a third may be very popular for certain businesses like Internet Service Providers (ISPs).

Dynamic routing protocols are different from each other in a number of ways, such as:

- [Classful versus classless routing protocols](#)
- [Interior versus exterior routing protocols](#)
- [Distance vector versus link-state protocols](#)

Classful versus classless routing protocols

Classful and classless routing refers to how the routing protocol handles the IP addresses. In classful addresses, there's the specific address and the host address of the server that address is connected to. Classless addresses use a combination of IP address and netmask.

Classless Inter-Domain Routing (CIDR) was introduced in 1993 (originally with [RFC 1519](#) and most recently with [RFC 4632](#)) to keep routing tables from getting too large. With classful routing, each IP address requires its own entry in the routing table. With classless routing, a series of addresses can be combined into one entry, potentially saving vast amounts of space in routing tables.

Current routing protocols that support classless routing, out of necessity, include RIPv2, BGP, IS-IS, and OSPF. Older protocols, such as RIPv1, do not support CIDR addresses.

Interior versus exterior routing protocols

The names interior and exterior are very descriptive. Interior routing protocols are designed for use within a contained network of limited size, whereas exterior routing protocols are designed to link multiple networks together. They can be used in combination in order to simplify network administration. For example, a network can be built with only border routers of a network running the exterior routing protocol, while all the routers on the network run the interior protocol. This prevents them from connecting outside the network without passing through the border. Exterior routers in such a configuration must have both exterior and interior protocols to communicate with the interior routers and outside the network.

Nearly all routing protocols are interior routing protocols. Only BGP is commonly used as an exterior routing protocol.

You may see interior gateway protocol (IGP) used to refer to interior routing protocols and exterior gateway protocol (EGP) used to refer to exterior routing protocols.

Distance vector versus link-state protocols

Every routing protocol determines the best route between two addresses using a different method. However, there are two main algorithms for determining the best route: distance vector and link-state.

Distance vector protocols

In distance vector protocols, routers are told about remote networks through neighboring routers. The distance part refers to the number of hops to the destination and, in more advanced routing protocols, these hops can be weighted by factors such as available bandwidth and delay. The vector part determines which router is the next step along the path for this route. This information is passed along from neighboring routers with routing update packets that keep the routing tables up to date. Using this method, an outage along a route is reported back along to the start of that route, ideally before the outage is encountered.

On distance vector protocols, [RFC 1058](#), which defines RIP v1, states the following:

Distance vector algorithms are based on the exchange of only a small amount of information. Each entity (gateway or host) that participates in the routing protocol is assumed to keep information about all of the destinations within the system. Generally, information about all entities connected to one network is summarized by a single entry, which describes the route to all destinations on that network.

There are four main weaknesses inherent in the distance vector method. Firstly, the routing information isn't discovered by the router itself, but is instead reported information that must be relied on to be accurate and up-to-date. The second weakness is that it can take a while for the information to make its way to all the routers who need the information; in other words, it can have slow convergence. The third weakness is the amount of overhead involved in passing these updates all the time. The number of updates between routers in a larger network can significantly reduce the available bandwidth. The fourth weakness is that distance vector protocols can end up with routing-loops. Routing loops are when packets are routed forever around a network, and often occur with slow convergence. The bandwidth required by these infinite loops will slow your network to a halt. There are methods of preventing these loops, however, so this weakness isn't as serious as it may first appear.

Link-state protocols

Link-state protocols are also known as shortest path first protocols. Where distance vector uses information passed along that may or may not be current and accurate, in link-state protocols each router passes along information only about the networks and devices that are directly connected to it. This results in a more accurate picture of the network topology around your router, allowing it to make better routing decisions. This information is passed between routers using link-state advertisements (LSAs). To reduce the overhead, LSAs are only sent out when information changes, compared to distance vector sending updates at regular intervals even if no information has changed. The more accurate network picture in link-state protocols greatly speed up convergence and avoid problems such as routing-loops.

Minimum configuration for dynamic routing

Dynamic routing protocols don't pay attention to routing updates from other sources, unless you specifically configure them to do so using CLI redistribute commands within each routing protocol.

The minimum configuration for any dynamic routing to function is to have dynamic routing configured on one interface on a FortiGate, and one other router configured as well. Some protocols require larger networks to function as designed.

Minimum configuration based on dynamic protocol

	BGP	RIP	OSPF / IS-IS
Interface	Yes	Yes	Yes
Network	Yes	Yes	Yes
AS	Local and neighbor	No	Yes
Neighbors	At least one	At least one	At least one
Version	No	Yes	No
Router ID	No	No	Yes

Comparison of dynamic routing protocols

Each dynamic routing protocol was designed to meet a specific routing need. Each protocol does some things well and other things not so well. For this reason, choosing the right dynamic routing protocol for your situation isn't an easy task.

Features of dynamic routing protocols

Each protocol is better suited for some situations over others.

Choosing the best dynamic routing protocol depends on the size of your network, speed of convergence required, the level of network maintenance resources available, what protocols the networks you connect to are using, and so on. For more information about these dynamic routing protocols, see ["RIP" on page 2110](#), ["BGP" on page 2187](#), ["OSPF" on page 2146](#), and ["IS-IS" on page 2224](#).

Comparing RIP, BGP, and OSPF dynamic routing protocols

Protocol	RIP	BGP	OSPF / IS-IS
Routing algorithm	Distance vector, basic	Distance vector, advanced	Link-state
Common uses	Small, non-complex networks	Network backbone, ties multinational offices together	Common in large, complex enterprise networks

Protocol	RIP	BGP	OSPF / IS-IS
Strengths	Fast and simple to implement	Graceful restart	Fast convergence
	Near universal support	BFD support	Robust
	Good when no redundant paths	Only needed on border routers	Little management overhead
		Summarize routes	No hop count limitation
			Scalable
Weakness	Frequent updates can flood network	Required full mesh in large networks can cause floods	Complex
	Slow convergence	Route flap	No support for unequal cost multipath routing
	Maximum 15 hops may limit network configuration	Load-balance multi-homed networks	Route summary can require network changes
		Not available on low-end routers	
Authentication	Optional authentication using text string or MD5 password. (RIP v1 has no authentication)		
IPv6 support	Only in RIPng	Only in BGP4+	Only in OSPF6 / Integrated IS-IS

Routing protocols

- Routing Information Protocol (RIP) uses classful routing, as well as incorporating various methods to stop incorrect route information from propagating, such as the poisoned horizon method. However, on larger networks its frequent updates can flood the network and its slow convergence can be a problem.
- Border Gateway Protocol (BGP) has been the core Internet backbone routing protocol since the mid-1990s, and is the most used interior gateway protocol (IGP). However, some configurations require full mesh connections which flood the network, and there can be route flap and load balancing issues for multihomed networks.
- Open Shortest Path First (OSPF) is commonly used in large enterprise networks. It is the protocol of choice, mainly due to its fast convergence. However, it can be complicated to setup properly.
- Intermediate System to Intermediate System (IS-IS) protocol allows routing of ISO's OSI protocol stack Connectionless Network Service (CLNS). IS-IS is an Interior Gateway Protocol (IGP) that's not intended to be used between Autonomous Systems (ASes). IS-IS is a link state protocol well-suited to smaller networks that's in widespread use and has near universal support on routing hardware.
- Multicast addressing is used to broadcast from one source to many destinations efficiently. Protocol Independent Multicast (PIM) is the protocol commonly used in enterprises, multimedia content delivery, and stock exchanges.

Routing algorithm

Each protocol uses a slightly different algorithm for choosing the best route between two addresses on the network. The algorithm is the intelligent part of a dynamic protocol because the algorithm is responsible for deciding which route is best and should be added to the local routing table. RIP and BGP use distance vector algorithms, where OSPF and IS-IS use link-state or a shortest path first algorithm.

Vector algorithms are essentially based on the number of hops between the originator and the destination in a route, possibly weighting hops based on how reliable, fast, and error-free they are.

The link-state algorithm used by OSPF and IS-IS is called the Dijkstra algorithm. Link-state treats each interface as a link and records information about the state of the interface. The Dijkstra algorithm creates trees to find the shortest paths to the routes it needs based on the total cost of the parts of the routes in the tree.

For more information about the routing algorithm used, see ["Distance vector versus link-state protocols" on page 2095](#).

Authentication

If an attacker gains access to your network, they can masquerade as a router on your network to either gain information about your network or disrupt network traffic. If you have a high quality firewall configured, it will help your network security and stop many of these types of threats. However, the main method for protecting your routing information is to use authentication in your routing protocol. Using authentication on a FortiGate and other routers, prevents access by attackers because all routers must authenticate with passwords, such as MD5 hash passwords, to ensure they are legitimate routers.

When you configure authentication on your network, ensure that you configure it the same way on all devices on the network. Failure to do so will create errors and outages as those forgotten devices fail to connect to the rest of the network.

For example, to configure an MD5 key of 123 on an OSPF interface called `ospf_test`, enter the following CLI commands:

```
config router ospf
  config ospf-interface
    edit ospf_test
      set authentication md5
      set md5-key 123
    next
  end
```

Convergence

Convergence is the ability of a networking protocol to re-route around network outages. Static routing can't do this. Dynamic routing protocols can all converge, but take various amounts of time to do this. Slow convergence can cause problems, such as network loops, which degrade network performance.

You may also hear robustness and redundancy used to describe networking protocols. In many ways, they're the same thing as convergence. Robustness is the ability to keep working even though there are problems, including configuration problems as well as network outages. Redundancy involves having duplicate parts that can continue to function in the event of some malfunction, error, or outage. It's relatively easy to configure dynamic routing protocols to have backup routers and configurations that will continue to function no matter what the network problem is, short of a total network failure.

IPv6 support

IPv4 addressing is in common use everywhere around the world. IPv6 has much larger addresses and it's used by many large companies and government departments. IPv6 isn't as common as IPv4 yet, but more companies are adopting it.

If your network uses IPv6, your dynamic routing protocol must support it. None of the dynamic routing protocols supported IPv6 originally, but they all have additions, expansions, or new versions that now support IPv6. For more information, see ["RIP" on page 2110](#), ["BGP" on page 2187](#), ["OSPF" on page 2146](#), and ["IS-IS" on page 2224](#).

When to adopt dynamic routing

Static routing is more than enough to meet your networking needs when you have a small network. However, as your network grows, the question you need to answer is at what point do you adopt dynamic routing in your networking plan and start using it in your network? The main factors in this decision are typically budget, current network size and topology, expected network growth, and available resources for ongoing maintenance.

Budget

When making any business decision, you must always consider your budget. Static routing doesn't involve special hardware, fancy software, or expensive training courses.

Dynamic routing can include all of these extra expenses. Any new hardware, such as routers and switches, need to support the routing protocols that you choose. Network management software and routing protocol drivers may also be necessary to help configure and maintain your more complex network. If the network administrators are not well versed in dynamic routing, you must budget either a training course or some hands-on learning time so they can administer the new network with confidence. Together, these factors can impact your budget.

Additionally, people will always account for network starting costs in the budgets but usually leave out the ongoing cost of network maintenance. Any budget must provide for the hours that will be spent on updating the network routing equipment and fixing any problems. Without that money in the budget, you may end up back at static routing before you know it.

Current network size and topology

As stated earlier, static routing works well on small networks. As those networks get larger, routing takes longer, routing tables get very large, and general performance isn't what it could be.

Topology is a concern as well. If all your computers are in one building, it is much easier to stay with static routing longer. However, connecting a number of locations will be easier with the move to dynamic routing.

If you have a network of 20 computers, you can still likely use static routing. If those computers are in two or three locations, static routing will still be a good choice for connecting them. Also, if you just connect to your ISP and don't worry about any special routing to do that, you're likely safe with just static routing.

If you have a network of 100 computers in one location, you can use static routing but it will be slower, more complex, and there won't be much room for expansion. If those 100 computers are spread across three or more locations, dynamic routing is the way to go.

If you have 1000 computers, you definitely need to use dynamic routing no matter how many locations you have.

Hopefully this section has given you an idea of what results you'll likely experience from different sized networks using different routing protocols. Your choice of which dynamic routing protocol to use is partly determined by the network size and topology.

Expected network growth

You may not be sure if your current network is ready for dynamic routing. However, if you're expecting rapid growth in the near future, it's a good idea to start planning for that growth now so you're ready for the coming expansion.

Static routing is very labor intensive. Each network device's routing table needs to be configured and maintained manually. If there's a large number of new computers being added to the network, they each need to have the static routing table configured and maintained. If devices are being moved around the network frequently, they must also be updated each time.

Instead, consider putting dynamic routing in place before the new computers are installed on the network. The installation issues can be worked out with a smaller and less complex network, and when the new computers or routers are added to the network there will be nowhere near the level of manual configuration required. Depending on the level of growth, the labor savings can be significant. For example, in an emergency you can drop a new router into a network or AS, wait for it to receive the routing updates from its neighbors, and then remove one of the neighbors. While the routes will not be the most effective possible, this method is much less work than static routing in the same situation, with less chance of mistakes.

Also, as your network grows and you add more routers, the new routers can help share the load in most dynamic routing configurations. For example, if you have 4 OSPF routers and 20,000 external routes, those few routers will be overwhelmed. But a network with 15 OSPF routers will be better able to handle that number of routes. However, be aware that adding more routers to your network will increase the amount of updates sent between the routers, which will use up a greater part of your bandwidth and use more bandwidth overall.

Available resources for ongoing maintenance

As explained in the budget section, there must be resources dedicated to ongoing network maintenance, upgrades, and troubleshooting. These resources include administrator hours to configure and maintain the network, training for the administrator (if needed), extra hardware and software as needed, and possible extra staff to help the administrator in emergencies. Without these resources, you'll quickly find the network reverting to static routing out of necessity. This is because:

- Routing software updates require time
- Routing hardware updates require time
- Office reorganizations or significant personnel movement require time from a networking point of view
- Networking problems that occur, such as failed hardware, require time to locate and fix the problem

If resources to accomplish these tasks are not budgeted, the tasks will either not happen at the required level to continue operation or not happen at all. This will result in both the network administration staff and the network users being very frustrated.

A lack of a maintenance budget will also result in an increasingly heavy reliance on static routing as the network administrators are forced to use quick fixes for problems that come up. This invariably involves going to static routing, and dropping the more complex and time-consuming dynamic routing.

Choosing a routing protocol

One of the hardest decisions in routing can be choosing which routing protocol to use on your network. It can be easy to decide when static routing won't meet your needs, but how can you tell which dynamic routing protocol is best for your network and situation?

Here's a brief look at the routing protocols, including their strongest and weakest points. The steps to choosing your routing protocol are:

1. [Answer questions about your network](#)
2. [Evaluate your chosen protocol](#)
3. [Implement your dynamic routing protocol](#)

Answer questions about your network

Before you can decide what is best for your situation, you need to examine the details of your situation, such as what your budget, equipment, and users are.

The following questions will help you form a clear idea of your routing needs:

How many computers or devices are on your network?

The number of computers or devices that you have on your network, and whether the devices are all in one location or distributed, matters. All routing protocols can be run on networks of any size, but it can be inefficient to run some routing protocols on very small networks. Also, routers and network hardware that support dynamic routing can be more expensive than more generic routers for static routing.

What applications typically run over the network?

Finding out what applications your users are running will help you determine their needs and the needs of the network regarding bandwidth, quality of service, and other such issues.

What level of service do the users expect from the network?

Different users have different expectations of the network. It's not critical for someone surfing the Internet to have 100% uptime, but it's required for a stock exchange network or a hospital.

Is there network expansion in your near future?

You may have a small network now, but if it will be growing quickly, you should plan for the expected size so you don't have to change technologies again down the road.

What routing protocols do your networks connect to?

This is most often how routing protocol decisions are made. You need to be able to communicate easily with your service provider and neighbors, so often people simply use what everyone else is using.

Is security a major concern?

Some routing protocols have levels of authentication and other security features built in, and others do not. If security is important to you, be aware of this.

What is your budget?

You need to know what both your initial and maintenance budget is. More robust and feature-laden routing protocols generally mean more resources are required to keep them working well. Also, more secure configurations require still more resources. This includes both set up costs and ongoing maintenance costs. If you ignore these costs, you risk having to drop the adoption of the new routing protocol mid-change.

Evaluate your chosen protocol

Once you've examined the features of the routing protocols listed above and chosen the one that best meets your needs, you can set up an evaluation or test installation of that protocol.

The test installation is generally set up in a sandbox configuration so it won't affect critical network traffic. The aim of the test installation is to prove that it will work on a larger scale on your network. You must ensure that the test installation mirrors your larger network well enough for you to discover any problems. If the test installation is too simple, these problems may not appear.

If your chosen protocol does not meet your goals, choose a different protocol and repeat the evaluation process until a protocol meets your needs or you change your criteria.

Implement your dynamic routing protocol

You've examined your needs, selected the best matching dynamic routing protocol, tested it, and now you're ready to implement it with confidence.

This guide will help you configure a FortiGate to support your chosen dynamic routing protocol. Refer to the various sections in this guide, as needed, during your implementation to help ensure a smooth transition. Examples for each protocol are included to show proper configurations for different types of networks.

Dynamic routing terminology

Dynamic routing is a complex subject. There are many routers on different networks and all of them can be configured differently. It's more complicated by the fact that each routing protocol has different names for similar features, as well as many features that you can configure for each protocol.

To better understand dynamic routing, the following sections provide explanations on common dynamic routing terms.

For more details about a term, as it applies to a dynamic routing protocol, see ["BGP" on page 2187](#), ["RIP" on page 2110](#), and ["OSPF" on page 2146](#).

Aggregated routes and addresses

Just as an aggregate interface combines multiple interfaces into one virtual interface, an aggregate route combines multiple routes into one route. This reduces the amount of space those routes require in the routing tables of the routers along that route. The trade-off is a small amount of processing to aggregate and de-aggregate the routes at either end.

The benefit of this method is that you can combine many addresses into one, potentially reducing the routing table size immensely. The weakness of this method is if there are holes in the address range you are aggregating, you need to decide if it is better to break it into multiple ranges, or accept the possibility of failed routes to the missing addresses.

For information about aggregated routes in BGP, see ["BGP" on page 2187](#).

To manually aggregate the range of IP addresses from 192.168.1.100 to 192.168.1.103:

1. Convert the addresses to binary:

```
192.168.1.100 = 11000000 10101000 00000001 01100100
192.168.1.101 = 11000000 10101000 00000001 01100101
192.168.1.102 = 11000000 10101000 00000001 01100110
192.168.1.103 = 11000000 10101000 00000001 01100111
```

2. Determine the maximum number of matching bits common to the addresses.

There are 30 bits in common, with only the last 2 bits being different.

3. Record the common part of the address:

```
11000000 10101000 00000001 0110010X = 192.168.1.100
```

4. For the netmask, assume all the bits in the netmask are 1, except those that are different (which are 0):

```
11111111 11111111 11111111 11111100 = 255.255.255.252
```

5. Combine the common address bits and the netmask:

```
192.168.1.100/255.255.255.252
```

Alternately, the IP mask may be written as a single number:

```
192.168.1.100/2
```

6. As required, set variables and attributes to declare that the routes have been aggregated, and which router did the aggregating.

Autonomous system

An Autonomous System (AS) is one or more connected networks that use the same routing protocol, and appear to be a single unit to any externally connected networks. For example, an ISP may have a number of customer networks connected to it, but to any networks connected externally to the ISP, it appears as one system or AS. An AS may also be referred to as a routing domain.

It should be noted that while OSPF routing takes place within one AS, the only part of OSPF that deals with the AS is the AS border router (ASBR).

There are multiple types of ASs, which are defined by how they are connected to other ASs. A multihomed AS is connected to at least two other ASs and has the benefit of redundancy. If one of those ASs goes down, your AS can still reach the Internet through its other connection. A stub AS has only one connection and can be useful in specific configurations where limited access is desirable.

Each AS has a number assigned to it, known as an ASN. In an internal network, you can assign any ASN you like (a private AS number), but for networks connected to the Internet (public AS), you need to have an officially registered ASN from the Internet Assigned Numbers Authority (IANA). ASNs from 1 to 64,511 are designated for public use.



NAs of January 2010, AS numbers are 4 bytes long, instead of the former 2 bytes. [RFC 4893](#) introduced 32-bit ASNs, which FortiGate support for BGP and OSPF.

Do you need your own AS?

The main factors in deciding if you need your own AS, or if you should be part of someone else's are:

- Exchanging external routing information
- Many prefixes should exist in one AS as long as they use the same routing policy
- When you use a different routing protocol than your border gateway peers. For example, your ISP uses BGP and you use OSPF.
- Connected to many other ASs (multihomed)

You shouldn't create an AS for each prefix on your network. You also shouldn't be forced into an AS just so someone else can make AS-based policy decisions on your traffic.

There can be only one AS for any prefix on the Internet. This is to prevent routing issues.

What AS number should you use?

In addition to overseeing IP address allocation and Domain Name Systems (DNS), the Internet Assigned Numbers Authority (IANA) assigns public AS numbers. The public AS numbers range from 1 to 64,511. The ASNs 0, 54272 to 64511, and 65535 are reserved by the IANA and shouldn't be used.

ASNs are assigned in blocks by the Internet Assigned Numbers Authority (IANA) to Regional Internet Registries (RIR), who then assign ASNs to companies within the geographic area of the RIR. These companies are usually ISPs, and to receive an ASN you must complete the application process of the local RIR and be approved before being assigned an ASN. The following table shows the names and regions of the RIRs:

AFRINIC	Serves the African continent
APNIC	Asia-Pacific, including China, India, and Japan
ARIN	American registry, including Canada and United States
LACNIC	Latin America, including Mexico, Caribbean, Central and South America
RIPE NCC	Europe, the Middle East, the former USSR, and parts of Central Asia

AS numbers from 64512 to 65534 are reserved for private use. Private AS numbers can be used for any internal networks with no outside connections to the Internet, such as test networks, classroom labs, and other internal-only networks that don't access the outside world. You can also configure border routers to filter out any private ASNs before routing traffic to the outside world. If you must use private ASNs with public networks, this is the only way to configure them. However, it's risky because many other private networks could be using the same ASNs and conflicts could happen. It would be like your local 192.168.0.0 network being made public and the resulting problems would be widespread.

In 1996, when [RFC 1930](#) was written, only 5,100 ASs had been allocated and a little under 600 ASs were actively routed in the global Internet. Since that time, many more public ASNs have been assigned, leaving only a small number. For this reason 32-bit ASNs (four-octet ASNs) were defined to provide more public ASNs. [RFC 4893](#) defines 32-bit ASNs, and a FortiGate supports these larger ASNs.

Area border router

Routers within an AS advertise updates internally and only to each other. However, routers on the edge of the AS must communicate both with routers inside their AS and routers external to their AS, which are often running a different routing protocol. These routers are called Area Border Routers (ABRs) or edge routers. ABRs often run multiple routing protocols in order to redistribute traffic between different ASs that are running different protocols, such as the edge between an ISP's IS-IS routing network and a large company's OSPF network.

OSPF defines ABRs differently from other routers. In OSPF, an ABR is an OSPF router that connects another AS to the backbone AS, and is a member of all the areas it connects to. An OSPF ABR maintains an LSA database for each area that it's connected to. The concept of the edge router is present, but it's the edge of the backbone instead of the edge of the OSPF supported ASs.

Neighbor routers

Routing involves routers communicating with each other. To do this, routers need to know information about each other. These routers are called neighbor routers and are configured in each routing protocol. Each neighbor has custom settings since some routers may have functionality that other routers lack. Neighbor routers are sometimes called peers.

Generally, neighbor routers must be configured and discovered by the rest of the network before they can be integrated into the routing calculations. This is a combination of the network administrator configuring the new router with its neighbor router addresses, and the routing network discovering the new router, such as the hello packets in OSPF. That discovery initiates communication between the new router and the rest of the network.

Route maps

Route maps are a way for a FortiGate to evaluate optimum routes for forwarding packets or suppressing the routing of packets to particular destinations. Compared to access lists, route maps support enhanced packet-matching criteria. In addition, route maps can be configured to permit or deny the addition of routes to the FortiGate routing table and make changes to routing information dynamically as defined through route-map rules.

Route maps can be used for limiting both received route updates and sent route updates. This can include the redistribution of routes learned from other types of routing. For example, if you don't want to advertise local static routes to external networks, you could use a route map to accomplish this.

The FortiGate compares the rules in a route map to the attributes of a route. The rules are examined in ascending order until one or more of the rules in the route map are found to match one or more of the route attributes.

As an administrator, route maps allow you to group a set of addresses together and assign them a meaningful name. During your configuration, you can use these route-maps to speed up configuration. The meaningful names also ensure that fewer mistakes are made during configuration.

The default rule in the route map, which the FortiGate applies last, denies all routes. For a route map to take effect, it must be called by a FortiGate routing process.

The syntax for route maps are:

```
config router route-map
  edit <route-map-name>
    set comments <comment>
    config rule
      edit <route-map-rule-id>
        set action {permit | deny}
        set match-*
        set set-*
        ...
      next
    next
  end
```


The `match-*` commands allow you to match various parts of a route. The `set-*` commands allow you to set routing information once a route is matched.

For an example of how route maps can be used to create receiving or sending “groups” in routing, see ["BGP" on page 2187](#).

Access lists

Use this command to add, edit, or delete access lists. Access lists are filters used by FortiGate routing processes. For an access list to take effect, it must be called by a FortiGate routing process (for example, a process that supports RIP or OSPF). Use `access-list6` for IPv6 routing.

Access lists can be used to filter which updates are passed between routers or which routes are redistributed to different networks and routing protocols. You can create lists of rules that will match all routes for a specific router or group of routers.

Each rule in an access list consists of a prefix (IP address and netmask), the action to take for this prefix (permit or deny), and whether to match the prefix exactly or match the prefix and a more specific prefix.



If you're setting a prefix of 128.0.0.0, use the format 128.0.0.0/1. The default route, 0.0.0.0/0 can't be exactly matched with an access-list. A prefix-list must be used for this purpose.

A FortiGate attempts to match a packet against the rules in an access list, starting at the top of the list. If it finds a match for the prefix, it takes the action specified for that prefix. If no match is found, the default action is deny.

The syntax for access lists is:

```
config router {access-list | access-list6}
  edit <access-list-name>
    set comments <comment>
    config rule
      edit <access-list-id>
        set action {permit | deny}
        set exact-match {enable | disable}
        set {prefix | prefix6} <prefix>
        set wildcard <wildcard>
      next
    next
  end
```

For an example of how access lists can be used to create receiving or sending “groups” in routing, see ["BGP" on page 2187](#).

BFD

Bidirectional Forwarding Detection (BFD) is a protocol that you can use to quickly locate hardware failures in the network. Routers running BFD send packets to each other at a negotiated rate. If packets from a BFD-protected router fail to arrive, that router is declared to be down. BFD communicates this information to the routing protocol and the routing information is updated.

BFD can run on an entire FortiGate, selected interfaces, or on a protocol, such as BGP, for all configured interfaces. The configuration hierarchy allows each lower level to override the BFD setting of the upper level. For example, if you enable BFD for an entire FortiGate, you can disable BFD for an interface or for BGP.

BFD neighbors establish if BFD is enabled in Open Shortest Path First (OSPF) or if BFD routers establish as neighbors.

The `config system` command allows you to configure whether BFD is enabled in a particular device or VDOM or individual interface, and how often the interface requires the sending and receiving of BFD information.

To configure BFD for an entire FortiGate - CLI:

```
config system settings
  set bfd {enable | disable}
  set bfd-desired-min-tx <ms>
  set bfd-required-min-rx <ms>
  set bfd-detect-mult <multiplier>
  set bfd-dont-enforce-src-port {enable | disable}
end
```

To configure BFD for an interface - CLI:

```
config system interface
  edit <interface-name>
    set bfd {global | enable | disable}
    set bfd-desired-min-tx <ms>
    set bfd-required-min-rx <ms>
    set bfd-detect-mult <multiplier>
  next
end
```

To show BFD neighbors - CLI:

```
get router {info | info6} bfd neighbor
```

To show BFD requests - CLI:

```
get router {info | info6} bfd requests
```

To configure BFD - CLI:

```
config router {bfd | bfd6}
  config neighbor
    edit <IP-address>
      set interface <interface-name>
    next
  end
```

BFD and static routes

BFD for static routes allows you to configure routing failover based on remote path failure detection. BFD removes a static route from the routing table if the FortiGate can't reach the route's destination and returns the route to the routing table if the route's destination is restored.

For example, you can add two static routes with BFD enabled. If one of the routes has a higher priority, all matching traffic uses that route. If BFD determines that the link to the gateway of the route with the higher priority is down, the higher priority route is removed from the routing table and all matching traffic uses the lower priority route. If the link to the gateway for the higher priority route comes back up, BFD adds the route back into the routing table and all matching traffic switches to use the higher priority route.

You can configure BFD for IPv4 and IPv6 static routes.

To configure BFD for static routes - CLI:

```
config router {static | static6}
  edit <sequence-number>
    set bfd {enable | disable}
  next
end
```

BFD and OSPF

You can configure BFD for Open Shortest Path First (OSPF) on a FortiGate. FortiGate supports BFD for OSPF for both IPv4 and IPv6.

To configure BFD for OSPF - CLI:

```
config router {ospf | ospf6}
  set bfd {enable | disable}
end
```

To enable BFD on a specific OSPF interface - CLI:

```
config router {ospf | ospf6}
  set bfd enable
  config {ospf-interface | ospf6-interface}
    edit <ID>
      set bfd {global | enable | disable}
    next
  end
end
```

BFD and BGP

While BGP can detect route failures, BFD can be configured to detect these failures more quickly, which allows for faster responses and improved convergence. This can be balanced with the bandwidth BFD uses in its frequent route checking.

The `config router bgp` commands allow you to set the addresses of the neighbor units that are also running BFD. Both units must be configured with BFD in order to use it.

To configure BFD for BGP - CLI:

```
config router bgp
  config neighbor
    edit <neighbor-IP-address>
      set bfd {enable | disable}
    next
  end
```

Controlling how routing changes affect active sessions

Dynamic routing changes can occur while a FortiGate is processing traffic. Routing changes that affect the routes being used for current sessions, may affect how the FortiGate continues to process the session. You can control

how active sessions are affected when dynamic routing changes occur that affects the routes the active sessions are using.

You can configure whether the FortiGate maintains the original routing for the sessions that are using the affected routes, or applies the routing table changes to the active sessions, which may cause destinations to change.

To configure how dynamic routing changes affect active sessions - CLI:

```
config system interface
  edit <interface_name>
    set preserve-session-route {enable | disable}
  next
```

where you set the following variables:

CLI option	Description
<interface_name>	The name of the interface where you want to configure how dynamic routing changes affect active sessions running through it.
enable (default)	All sessions passing through the interface when the routing changes occur, are allowed to finish and aren't affected by the routing changes.
disable	When a routing change occurs, the new routing table is applied to the active sessions passing through the interface. The routing changes may cause the destinations of the sessions to change.

IPv6 in dynamic routing

Unless otherwise stated, routing protocols apply to IPv4 addressing. This is the standard address format used. However, IPv6 is becoming more popular and new versions of the dynamic routing protocols have been introduced.

Dynamic routing supports IPv6 on a FortiGate. The new versions of these protocols and the corresponding RFCs are:

- **RIP next generation (RIPng)** — [RFC 2080](#) - Routing Information Protocol next generation (RIPng). See RIP and IPv6.
- **BGP4+** — [RFC 2545](#), and [RFC 2858](#) Multiprotocol Extensions for IPv6 Inter-Domain Routing, and Multiprotocol Extensions for BGP-4 (MP-BGP) respectively. See BGP and IPv6.
- **OSPFv3** — [RFC 2740](#) Open Shortest Path First version 3 (OSPFv3) for IPv6 support. See OSPFv3 and IPv6.
- **Integrated IS-IS** — [RFC 5308](#) for IPv6 support. See Integrated IS-IS.

As with most advanced routing features on a FortiGate, IPv6 settings for dynamic routing protocols must be enabled before they will be visible in the GUI. To enable IPv6 configuration in the GUI, enable it in **System > Admin > Settings**. Alternatively, you can directly configure IPv6 for RIP, BGP, or OSPF protocols using CLI commands.

RIP

Routing Information Protocol (RIP) is a distance-vector routing protocol intended for small, relatively homogeneous networks. Its widespread use started when an early version of RIP was included with BSD v4.3 Linux as the routed daemon. The Bellman–Ford algorithm, which is the routing algorithm used by RIP, first saw widespread use as the initial routing algorithm of the ARPANET.

RIP has many benefits. It is well suited to smaller networks, has near universal support on routing hardware, is quick to configure, works well if there are no redundant paths, and is in widespread use. However, because RIP updates are sent out node-by-node, it can be slow to find a path around network outages. RIP also lacks good authentication, cannot choose routes based on different quality of service methods, and can create network loops if you are not careful.

The FortiGate implementation of RIP supports RIP version 1 (see [RFC 1058](#)), RIP version 2 (see [RFC 2453](#)), and the IPv6 version RIPng (see [RFC 2080](#)).

RIPv1

In 1988, RIP version 1 (RIPv1) was released. It is defined in [RFC 1058](#). It uses classful addressing and uses broadcasting to send out updates to router neighbors. There's no subnet information included in the routing updates in classful routing. It doesn't support CIDR addressing and subnets must all be the same size. Also, route summarization isn't possible. RIPv1 has no router authentication method, so it's vulnerable to attacks through packet sniffing and spoofing.

RIPv2

In 1993, RIP version 2 (RIPv2) was developed to deal with the limitations of RIPv1. It wasn't standardized until 1998. This new version supports classless routing and subnets of various sizes. Router authentication was added, which supports MD5. MD5 hashes are an older encryption method, but this is much improved over no security at all. In RIPv2, the hop count limit remained at 15, in order to be backwards compatible with RIPv1. It also uses multicasting to send the entire routing table to router neighbors, which reduces the traffic for devices that aren't participating in RIP routing. Routing tags were also added, which allow internal routes or redistributed routes to be identified as such.

RIPng

RIPng, defined in [RFC 2080](#), is an extension of RIPv2 and is designed to support IPv6. However, RIPng varies from RIPv2 in that it's not fully backwards compatible with RIPv1. RIPng doesn't support RIPv1 update authentication and relies on IPsec instead. It doesn't allow the attaching of tags to routes, as in RIPv2. RIPng requires specific encoding of the next hop for a set of route entries, unlike RIPv2 that encodes the next-hop into each route entry.

RIP terminology and parts

Before you can understand how RIP functions, you need to understand some of the main concepts and parts of RIP.

RIP and IPv6

RIP Next Generation (RIPng) is a new version of RIP and includes support for IPv6.

The FortiGate `config router ripng` command is almost the same as the `config router rip` command, except that IPv6 addresses are used. Also, if you're going to use prefix or access lists with RIPng, you must use the `config router access-list6` or `config prefix-list6` versions of those commands.

If you want to troubleshoot RIPng, it's the same as with RIP but specify the different protocol and use IPv6 addresses. This applies to commands such as `get router info6` when you want to see the routing table or other related information.

If you want to route IPv4 traffic over an IPv6 network, you can use the command `config system ipv6-tunnel` to configure the FortiGate to do this. The IPv6 interface is configured under `config system interface`. All subnets between the source and destination addresses must support IPv6. This command isn't supported in transparent mode.

For example, if you want to set up a tunnel on the port1 interface starting at 2002:C0A8:3201:: on your local network and tunnel it to address 2002:A0A:A01::, where it will need access to an IPv4 network again, use the following CLI commands:

```
config system ipv6-tunnel
  edit test_tunnel
    set destination 2002:A0A:A01::
    set interface port1
    set source 2002:C0A8:3201::
  next
```

The CLI commands associated with RIPng include:

```
config router ripng
config router access-list6
config router prefix-list6
config system ipv6-tunnel
get router info6 *
```

Default information originate option

The default information originate option is the second advanced option for RIP in the FortiGate GUI, right after metric. Enabling default-information-originate will generate and advertise a default route into the FortiGate device's RIP-enabled networks. The generated route may be based on routes learned through a dynamic routing protocol, routes in the routing table, or both. RIP does not create the default route unless you use the always option.

Select **Disable** if you experience any issues or if you wish to advertise your own static routes into RIP updates.

You can enable or disable default-information-originate in **Router > Dynamic > RIP**, under **Advanced Options**, or use the CLI.

The CLI commands associated with default information originate include:

```
config router rip
  set default-information-originate {enable | disable}
end
```

Update, timeout, and garbage timers

RIP uses various timers to regulate its performance including an update timer, a timeout timer, and a garbage timer. The FortiGate default timer settings (30, 180, and 120 seconds) are effective in most configurations. If you change these settings, ensure that the new settings are compatible with local routers and access servers.



The timeout period should be at least three times longer than the update period. If the update timer is smaller than the timeout or garbage timers, you'll experience an error.

You can set the three RIP timers in **Router > Dynamic > RIP**, under **Advanced Options**, or use the CLI.

To configure garbage, timeout, and update timers - CLI:

```
config router rip
  set timeout-timer <seconds>
  set update-timer <seconds>
  set garbage-timer <seconds>
end
```

Update timer

The update timer determines the interval between routing updates. This value is usually set to 30 seconds. There's some randomness added to help prevent network traffic congestion, which could result from all routers attempting to update their neighbors simultaneously. The update timer should be at least three times smaller than the timeout timer or you'll experience an error.

If you're experiencing significant RIP traffic on your network, you can increase this interval to send fewer updates per minute. However, ensure you increase the interval for all the routers on your network or you'll experience timeouts that will degrade your network speed.

Timeout timer

The timeout timer is the maximum amount of time (in seconds) that a route is considered reachable while no updates are received for the route. This is the maximum time a FortiGate will keep a reachable route in the routing table while no updates for that route are received. If the FortiGate receives an update for the route before the timeout period expires, the timer is restarted. The timeout period should be at least three times longer than the update period or you'll experience an error.

If you're experiencing problems with routers not responding in time to updates, increase this timer. However, remember that longer timeout intervals result in longer overall update periods. It may be a considerable amount of time before the FortiGate is done waiting for all the timers to expire on unresponsive routes.

Garbage timer

The garbage timer is the amount of time (in seconds) that a FortiGate advertises a route as being unreachable before deleting the route from the routing table. If this timer is shorter, it will keep more up-to-date routes in the routing table and remove older ones faster. This will result in a smaller routing table, which is useful if you have a very large network, or if your network changes frequently.

Authentication and key chain

RIP version 2 (RIPv2) uses authentication keys to ensure that the routing information exchanged between routers is reliable. RIP version 1 (RIPv1) has no authentication. For authentication to work, both the sending and receiving routers must be set to use authentication and must be configured with the same keys.

The sending and receiving routers need to have their system dates and times synchronized to ensure both ends are using the same keys at the proper times. However, you can overlap the key lifetimes to ensure that a key is always available even if there's some difference in the system times.

A key chain is a list of one or more authentication keys, including the send and receive lifetimes for each key. Keys are used for authenticating routing packets only during the specified lifetimes. A FortiGate migrates from one key to the next according to the scheduled send and receive lifetimes.

The key-chain command is a CLI router command. You use this command to manage RIPv2 authentication keys. You can add, edit, or delete keys identified by the specified key number.

This example shows how to configure a key chain with two keys that are valid sequentially in time. This example creates a key chain called `rip_key` that has a password of "fortinet". The accepted and send lifetimes are both set to the same values: a start time of 9:00 am on February 23, 2010 and an end time of 9:00 am on March 17, 2010. A second key is configured with a password of "my_fortigate" that is valid from March 17, 2010 9:01am to April 1 2010 9:00am. This "rip_key" key chain is then used on the port1 interface in RIP.

```
config router key-chain
  edit rip_key
    config key
      edit 1
        set accept-lifetime 09:00:00 23 02 2010 09:00:00 17 03 2010
        set key-string "fortinet"
        set send-lifetime 09:00:00 23 02 2010 09:00:00 17 03 2010
      next
      edit 2
        set accept-lifetime 09:01:00 17 03 2010 09:00:00 1 04 2010
        set key-string "my_fortigate"
        set send-lifetime 09:01:00 17 03 2010 09:00:00 1 04 2010
      next
    next
  end
config router rip
  config interface
    edit port1
      set auth-keychain "rip_key"
    next
  end
```

The CLI commands associated with authentication keys include:

```
config router key-chain

config router rip
  config interface
    edit <interface>
      set auth-keychain <key_chain_name>
      set auth-mode {none | text | md5}
      set auth-string <password_string>
    next
  end
```

Access lists

Access lists are filters used by the FortiGate device's RIP and OSPF routing. An access list provides a list of IP addresses and the action to take for them. Essentially, an access list makes it easy to group addresses that will

be treated the same way into the same group, independent of their subnets or other matching qualities. You add a rule for each address or subnet that you want to include and specify the action to take for it. For example, if you want all traffic from one department to be routed a particular way, even in different buildings, you can add all of the addresses to an access list and then handle that list all at once.

Each rule in an access list consists of a prefix (IP address and netmask), the action to take for this prefix (permit or deny), and whether to match the prefix exactly or to match the prefix and any more specific prefix.

A FortiGate attempts to match a packet against the rules in an access list, starting at the top of the list. If it finds a match for the prefix, it takes the action specified for that prefix. If a match isn't found, the default action is deny.

Access lists greatly speed up configuration and network management. When there's a problem, you can check each list instead of individual addresses. Also, it's easier to troubleshoot because if all addresses on one list have problems, many possible causes can be eliminated right away.

If you're using the RIPng or OSPF+ IPv6 protocols, you'll need to use access-list6, which is the IPv6 version of the access list. The only difference is that access-list6 uses IPv6 addresses.

For example, if you want to create an access list called `test_list` that only allows an exact match of 10.10.10.10 and 11.11.11.11, enter the command:

```
config router access-list
  edit test_list
    config rule
      edit 1
        set prefix 10.10.10.10 255.255.255.255
        set action permit
        set exact-match enable
      next
      edit 2
        set prefix 11.11.11.11 255.255.255.255
        set action permit
        set exact-match enable
      next
    next
  end
```

Another example is if you want to deny ranges of addresses in IPv6 that start with the IPv6 equivalents of 10.10.10.10 and 11.11.11.11, enter the access-list6 command, as follows:

```
config router access-list6
  edit test_list_ip6
    config rule
      edit 1
        set prefix6 2002:A0A:A0A:0:0:0:0:0/48
        set action deny
      next
      edit 2
        set prefix6 2002:B0B:B0B:0:0:0:0:0/48
        set action deny
      next
    next
  end
```

To use an access list, you must call it from a routing protocol, such as RIP. The following example uses the access list from the previous example, called `test_list`, to match routes coming in on the port1 interface. When

there's a match, it will add 3 to the hop count metric for those routes to artificially increase. Enter the following command:

```
config router rip
  config offset-list
    edit 5
      set access-list test_list
      set direction in
      set interface port1
      set offset 3
      set status enable
    next
  end
```

If you are setting a prefix of 128.0.0.0, use the format 128.0.0.0/1. The default route, 0.0.0.0/0 can't be exactly matched with an access list. A prefix list must be used for this purpose

How RIP works

As one of the original modern dynamic routing protocols, RIP is straightforward. Its routing algorithm isn't complex and there are some options that allow fine tuning. It is relatively simple to configure RIP on a FortiGate.

From [RFC 1058](#):

Distance vector algorithms are based on the exchange of only a small amount of information. Each entity (gateway or host) that participates in the routing protocol is assumed to keep information about all of the destinations within the system. Generally, information about all entities connected to one network is summarized by a single entry, which describes the route to all destinations on that network.

RIP versus static routing

RIP was one of the earliest dynamic routing protocols to work with IP addresses. As such, it's not as complex as more recent protocols. However, RIP is a big step forward from simple static routing.

While RIP may be slow in response to network outages, static routing has zero response. The same is true for convergence; static routing has zero convergence. Both RIP and static routing have a limited hop count, so it's not a strength or a weakness. Count to infinity can be a problem, but can typically be fixed as it happens, or is the result of a network outage that would cause even worse problems on a static routing network.

This compares to static routing where each time a packet needs to be routed, the FortiGate can send it only to the next hop towards the destination. That next hop then forwards it, and so on until it arrives at its destination. RIP keeps more routing information on each router so the FortiGate can send the packet further towards its destination before it has to be routed again toward its destination. RIP uses a smaller amount of table lookups, and therefore fewer network resources, than static routing. Also, since RIP is updated on neighboring routes, it's aware of new routes or dead routes that static routing would not be aware of.

Overall, RIP is a large step forward when compared to static routing.

RIP hop count

RIP uses hop count as the metric for choosing the best route. A hop count of 1 represents a network that is connected directly to a FortiGate, while a hop count of 16 represents a network that can't be reached. Each network that a packet travels through to reach its destination usually counts as one hop. When the FortiGate

compares two routes to the same destination, it adds the route having the lowest hop count to the routing table. As you can see in ["RIP packet structure" on page 2120](#), the hop count is part of a RIP v2 packet.

Similarly, when RIP is enabled on an interface, the FortiGate sends RIP responses to neighboring routers on a regular basis. The updates provide information about the routes in the FortiGate device's routing table, subject to the rules that you specify for advertising those routes. You can specify how often the FortiGate sends updates, the period of time a route can be kept in the routing table without being updated, and for routes that aren't updated regularly, you can specify the period of time that the unit advertises a route as unreachable before it's removed from the routing table.

If hops are weighted higher than one, it's very easy to reach the upper limit. This higher weighting effectively limits the size of your network, depending on the numbers used. Merely changing from the default of 1.0 to 1.5 will lower the effective hop count from 15 to 10. This is acceptable for smaller networks, but can be a problem as your network expands over time.

In RIP, you can use the `offset` command to artificially increase the hop count of a route. Doing this will make this route less preferred and, in turn, it will get less traffic. Offsetting routes is useful when you have network connections that have different bandwidths, levels of reliability, or costs. In each of these situations you still want the redundancy of multiple route access, but you don't want the bulk of your traffic using these less preferred routes. For an example of RIP offset, see ["Access lists" on page 2113](#).

The Bellman–Ford routing algorithm

The routing algorithm used by RIP was first used in 1967 as the initial routing algorithm of the ARPANET. The Bellman–Ford algorithm is distributed because it involves a number of nodes (routers) within an Autonomous system, and consists of the following steps:

1. Each node calculates the distances between itself and all other nodes within the AS and stores this information as a table.
2. Each node sends its table to all neighboring nodes.
3. When a node receives distance tables from its neighbors, it calculates the shortest routes to all other nodes and updates its own table to reflect any changes.

To examine how this algorithm functions let us look at a network with four routers: routers 1 through 4. The distance from Router1 to Router2 is 2 hops, Router1 to Router3 is 3 hops, and Router2 to Router3 is 4 hops. Router4 is only connected to Router2 and Router3, each distance being 2 hops.

1. Router1 finds all of the distances to the other three routers: Router 2 is 2, Router 3 is 3. Router1 doesn't have a route to Router4.
2. Router2, Router3, and Router4 perform the same calculations from their point of views.
3. Once Router1 gets an update from Router 2 or Router3, it will get their route to Router4. At that point, it now has a route to Router4 and installs that in its local table.
4. If Router1 gets an update from Router3 first, it has a hop count of 5 to reach Router4, but when Router2 sends its update, Router1 will go with Router2's shorter 4 hops to reach Router4. Future updates don't change this unless they are shorter than 4 hops or the routing table route goes down.

Example of a RIP algorithm


Step 1

Router1 finds the distance to other routers in the network.

It currently has no route to Router4.

Router1 routing table:


- Distance to Router2 = 2 hops
- Distance to Router3 = 3 hops



Step 2

All routers do the same as Router1 and send out updates containing their routing table.

Note that Router1 and Router4 don't update each other, but rely on Router2 and Router3 to pass along accurate updates.




Step 3

Each router looks at the updates it has received and adds any new or shorter routes to its table.

Router1's updated table:

- Distance to Router2 = 2 hops
- Distance to Router3 = 3 hops

- Distance to Router4 = 4 or 5 hops




Step 4

Router1 installs the shortest route to Router4 and the other routes to it are removed from the routing table.

Router1's complete table:

- Distance to Router2 = 2 hops
- Distance to Router3 = 3 hops
- Distance to Router4 = 4 hops



The good part about the Bellman-Ford algorithm in RIP is that the router only uses the information it needs from the update. If there are no newer, better routes than the ones the router already has in its routing table, there's no need to change its routing table. And no change means no additional update, and therefore less traffic. But even when there's update traffic, the RIP packets are very small so it takes many updates to affect overall network bandwidth. For more information about RIP packets, see ["RIP packet structure" on page 2120](#).

The main disadvantage of the Bellman–Ford algorithm in RIP is that it doesn't take weightings into consideration. While it's possible to assign different weights to routes in RIP, doing so severely limits the effective network size by reducing the hop count limit. Also, other dynamic routing protocols can take route qualities, such as reliability or delay, into consideration to provide not only the physically shortest routes but also the fastest or more reliable routes.

Another disadvantage of the Bellman-Ford algorithm is due to the slow updates passed from one RIP router to the next. This results in a slow response to changes in the network topology, which in turn results in more attempts to use routes that are down and that wastes time and network resources.

Passive versus active RIP interfaces

Normally, the FortiGate routing table is kept up to date by periodically asking the neighbors for routes, and sending your routing updates out. This has the downside of generating a lot of extra traffic for large networks. The solution to this problem is passive interfaces.

A standard interface that supports RIP is active, by default. It sends and receives updates by actively communicating with its neighbors. A passive RIP interface doesn't send out updates. It only listens to the updates of other routers. This is useful in reducing network traffic, and if there are redundant routers in the network that will send out essentially the same updates all the time.

The following example shows how to create a passive RIPv2 interface on port1 using MD5 authentication and a key chain called `passiveRIPv2`, which has already been configured. Note that in the CLI, you enable passive by disabling `send-version2-broadcast`.

To create a passive RIP interface - GUI:

1. Go to **Router > Dynamic > RIP**.
2. Next to **Interfaces**, select **Create**.
3. Select port1 as the **Interface**.
4. Select 2 as both the **Send Version** and **Receive Version**.
5. Select MD5 for **Authentication**.
6. Select the `passiveRIPv2` **Key-chain**.
7. Select **Passive Interface**.
8. Select **OK** to accept this configuration and return to the main RIP display page.

To create a passive RIP v2 interface on port1 using MD5 authentication - CLI:

```
config router rip
  config interface
    edit port1
      set send-version2-broadcast disable
      set auth-keychain "passiveRIPv2"
      set auth-mode md5
      set receive-version 2
      set send-version 2
    next
  end
```

RIP packet structure

It's hard to fully understand a routing protocol without knowing what information is carried in its packets. Knowing what information is exchanged between routers and how it's exchanged will help you to better understand the RIP protocol and better configure your network for it.

This section provides information about the contents of RIPv1 and RIPv2 packets.

RIP version 1

RIP version 1 (RIPv1), or RIP IP, packets are 24 bytes in length with some empty areas left for future expansion.

RIP IP packets

1-byte command	1-byte version	2-byte zero field	2-byte AFI	2-byte zero field
4-byte IP address	4-byte zero field	4-byte zero field	4-byte metric	

A RIPv1 packet contains the following fields:

- **Command:** Indicates whether the packet is a request or a response. The request asks that a router send all or part of its routing table. The response can be an unsolicited regular routing update or a reply to a request. Responses contain routing table entries. Multiple RIP packets are used to convey information from large routing tables.
- **Version:** Specifies the RIP version used. This field can signal different, potentially incompatible versions.
- **Zero field:** This field defaults to zero and is not used by [RFC 1058](#) RIP.
- **Address-family identifier (AFI):** Specifies the address family used. RIP is designed to carry routing information for several different protocols. Each entry has an address-family identifier to indicate the type of address being specified. The AFI for IP is 2.
- **IP Address:** Specifies the IP address for the entry.
- **Metric:** This is the number of hops or routers traversed along the route on its trip to the destination. The metric is between 1 and 15 for that number of hops. If the route is unreachable, the metric is 16.

RIP version 2

RIP version 2 (RIPv2) has more features than RIPv1, which is reflected in its packets that carry more information. All but one of the empty zero fields in RIPv1 packets are used in RIPv2.

RIPv2 packets

1-byte command	1-byte version	2-byte unused	2-byte AFI	2-byte route tag
4-byte IP address	4-byte subnet	4-byte next hop	4-byte metric	

A RIPv2 packet contains the fields described above for RIPv1, as well as the following:

- **Unused:** Has a value set to zero and is intended for future use
- **Route tag:** Provides a method for distinguishing between internal routes learned by RIP and external routes learned from other protocols.

- **Subnet mask:** Contains the subnet mask for the entry. If this field is zero, no subnet mask has been specified for the entry.
- **Next hop:** Indicates the IP address of the next hop to which packets for the entry should be forwarded.

Troubleshooting RIP

This section is about troubleshooting RIP. For general troubleshooting information, see the *Troubleshooting Handbook*.

Routing loops

Normally in routing, a path between two addresses is chosen and traffic is routed along that path from one address to the other. When there's a routing loop, that normal path doubles back on itself, creating a loop. When there are loops, the network has problems getting information to its destination. Loops also prevent the network from returning to the source to report the inaccessible destination.

A routing loop occurs when a normally functioning network has an outage and one or more routers are offline. When packets encounter this, they attempt an alternate route maneuver around the outage. During this phase, it's possible for a route to be attempted that involves going back a hop, and trying a different hop forward. If that hop forward is also blocked by the outage, a hop back, and possibly the original hop forward, may be selected. If this continues, it can consume not only network bandwidth but also many resources on the affected routers. The worst part is this situation will continue until the network administrator changes the router settings or the downed routers come back online.

Effect of routing loops on the network

In addition to this "traffic jam" of routed packets, every time the routing table for a router changes, that router sends an update out to all of the RIP routers connected to it. In a network loop, it's possible for a router to change its routes very quickly as it tries and fails along these new routes. This can quickly result in a flood of updates being sent out, which can effectively grind the network to a halt until the problem is fixed.

How to spot a routing loop

Anytime network traffic slows down, you'll ask yourself if it's a network loop. Slowdowns are often normal, aren't a full stoppage, and normal traffic resumes in a short period of time.

If the slow down is a full halt of traffic or a major slowdown that doesn't return to normal quickly, you need to do serious troubleshooting quickly.

If you're not running SNMP, link health monitoring, or you have non-Fortinet routers in your network, you can use networking tools, such as ping and traceroute, to define the outage on your network and begin to fix it. Ping, traceroute, and other basic troubleshooting tools are largely the same between static and dynamic and are covered in ["Advanced static routing" on page 2062](#).

Check your logs

If your routers log events to a central location, it can be easy to check your network logs for any outages.

In the FortiGate GUI, go to **Log & Report**. You should look at both event logs and traffic logs. Events to look for generally fall under CPU and memory usage, interfaces going offline (due to link health monitoring), and other similar system events.

Once you have found and fixed your network problem, you can go back to the logs and create a report to better see how things developed during the problem. This type of forensics analysis can better help you prepare for next time.

Use SNMP network monitoring

If your network had no problems one minute and slows to a halt the next, chances are something changed to cause that problem. Most of the time an offline router is the cause and once you find that router and bring it back online, things will return to normal.

If you can enable a hardware monitoring system such as SNMP or sFlow on your routers, you can be notified of the outage and its location as soon as it happens.

Ideally, you can configure SNMP on all FortiGate routers and be alerted to all outages as they occur.

To use SNMP to detect potential routing loops - GUI:

1. Go to **System > Config > SNMP**.
2. Enable **SMTP Agent** and select **Apply**.
Optionally, enter the **Description**, **Location**, and **Contact** information for this device for easier location of the problem report.
3. Under **SNMP v1/v2** or **SNMP v3** as appropriate, select **Create New**.

SNMP v3

User Name	Enter the SNMP user ID.
Security Level	Select authentication or privacy as desired. Select the authentication or privacy algorithms to use and enter the required passwords.
Notification Host	Enter the IP addresses of up to 16 hosts to notify.
Enable Query	Select. The Port should be 161. Ensure that your security policies allow ports 161 and 162 (SNMP queries and traps) to pass.

SNMP v1/v2

Hosts	Enter the IP addresses of up to 8 hosts to notify.
Queries	Enable v1 and/or v2 as needed. The Port should be 161. Ensure that your security policies allow port 161 to pass.
Traps	Enable v1 and/or v2 as needed. The Port should be 162. Ensure that your security policies allow port 162 to pass.

4. Select the events for which you want notification. For routing loops this should include **CPU usage is high**, **Memory is low**, and possibly **Log disk space is low**. If there are problems the log will fill up quickly and the FortiGate device's resources will be overused.
5. Configure SNMP host (manager) software on your administration computer. This will monitor the SNMP information sent out by the FortiGate. Typically you can configure this software to alert you to outages or CPU spikes that may indicate a routing loop.

Use link health monitoring

Another tool available to you on a FortiGate is the link health monitor. You can detect possible routing loops with link health monitors. You can configure the FortiGate to ping a gateway at regular intervals to ensure it's online and working. When the gateway isn't accessible, that interface is marked as down.

For more information about link health monitoring, see ["Link health monitor" on page 2025](#).

Use email alerts for failed gateways

You can detect possible routing loops with email alerts.

To configure notification of failed gateways - GUI:

1. Go to **Log & Report > Report > Local** and enable **Email Generated Reports**.
2. Enter your email details.
3. Select **Apply**.

You might also want to log CPU and Memory usage because a network outage will cause your CPU activity to spike.



If you have VDOMs configured, you will have to enter the basic SMTP server information in the Global section, and the rest of the configuration within the VDOM that includes this interface.

After this configuration, when this interface on the FortiGate can't connect to the next router, the FortiGate brings down the interface and alert you with an email about the outage.

Look at the packet flow

If you want to see what is happening on your network, look at the packets traveling on the network. This is the same idea as police pulling over a car and asking the driver where they have been and what the conditions were like.

The method used in the troubleshooting sections ["Debugging IPv6 on RIPng" on page 2124](#) and on debugging the packet flow also apply here. In this situation, you're looking for routes that have metrics higher than 15, as this indicates they are unreachable.

Ideally, if you debug the flow of the packets and record the routes that are unreachable, you can create an accurate picture of the network outage.

Action to take on discovering a routing loop

Once you've mapped the problem on your network and determined that it's a routing loop, there are a number of steps you can take to correct it:

1. Get any offline routers back online. This may be a simple reboot or you may have to replace hardware. Often, this first step will restore your network to its normal operation once the routing tables are updated.
2. Change your routing configuration on the edges of the outage. Even if step 1 brought your network back online, you should consider making changes to improve your network before the next outage occurs. These changes can include configuring features like holddowns and triggers for updates, split horizon, and poison reverse updates.

Holddowns and triggers for updates

One of the potential problems with RIP is the frequent routing table updates that are sent every time there's a change to the routing table. If your network has many RIP routers, these updates can start to slow your network down. Also, if you have a particular route that has bad hardware, it might be going up and down frequently, which will generate an overload of routing table updates.

One of the most common solutions to this problem is to use holddown timers and triggers for updates. These slow down the updates that are sent out and help prevent a potential flood.

Holddown timers

The holddown timer activates when a route is marked down. Until the timer expires, the router doesn't accept any new information about that route. This is very useful if you have a flapping route because it'll prevent your router from sending out updates and being part of the problem in flooding the network. The potential downside is if the route comes back up before the timer expires, that route will be unavailable for that period of time. This is only a problem if this is a major route used by the majority of your traffic. Otherwise, this is a minor problem as traffic can be re-routed around the outage.

Triggers

Triggered RIP is an alternate update structure that is based around limiting updates to only specific circumstances. The most basic difference is that the routing table is only updated when a specific request is sent to update, instead of every time the routing table changes. Updates are also triggered when a unit is powered on, which can include the addition of new interfaces or devices to the routing structure, or devices returning to being available after being unreachable.

Split horizon and poison reverse updates

Split horizon is best explained with an example. If there are three routers linked serially, called routerA, routerB, and routerC. RouterA is only linked to routerB, RouterC is only linked to routerB, and routerB is linked to both routerA and routerC. To get to routerC, routerA must go through routerB. If the link to routerC goes down, it's possible that routerB will try to use routerA's route to get to routerC. This route is A-B-C, so it'll loop endlessly between routerA and routerB.

This situation is called a split horizon because from routerB's point of view the horizon stretches out in each direction but in reality it's only on one side. Poison reverse is the method used to prevent routes from running into split horizon problems. Poison reverse "poisons" routes away from the destination that use the current router in their route to the destination. This poisoned route is marked as unreachable for routers that can't use it. In RIP, this means that the route is marked with a distance of 16.

Debugging IPv6 on RIPng

The debug commands are very useful to see what is happening on the network at the packet level. There are a few changes to debugging the packet flow when debugging IPv6.

The following CLI commands specify both IPv6 and RIP, so only RIPng packets will be reported. The output from these commands will show you the RIPng traffic on your FortiGate unit, including RECV, SEND, and UPDATE actions.

The addresses are in IPv6 format.

```
diagnose debug enable
diagnose ipv6 router rip level info
```

```
diagnose ipv6 router rip all enable
```

These three commands:

- Turn on debugging, in general
- Set the debug level to information, which is a verbose reporting level
- Turn on all RIP router settings

Part of the information displayed from the debugging is the metric (hop count). If the metric is 16, that destination is unreachable since the maximum hop count is 15.

In general, you should see an update announcement, followed by the routing table being sent out, and a reply received in response.

For more information, see ["Testing the IPv6 RIPng information" on page 2146](#).

Simple RIP example

This is an example of a typical medium-sized network configuration using RIP routing.

Your company has 3 small local networks, one for each department. These networks are connected by RIP, and then connected to the Internet. Each subnet has more than one route for redundancy. There are two central routers that are both connected to the Internet and to the other networks. If one of those routers goes down, the whole network can continue to function normally.

The ISP is running RIP, so no importing or exporting routes is required on the side of the network. However, since the internal networks have static networking running, those will need to be redistributed through the RIP network.

To keep the example simple, there will be no authentication of router traffic.

With RIP properly configured, if the device fails or temporarily goes offline, the routes will change and traffic will continue to flow. RIP is good for a smaller network due to its lack of complex configurations.

Network layout and assumptions

Basic network layout

Your company has 3 departments each with their own network: Sales, R&D, and Accounting. Each network has routers that are not running RIP and FortiGate devices running RIP.

The R&D network has two RIP routers, and each is connected to both other departments as well as being connected to the Internet through the ISP router. The links to the Internet are indicated in black.

The three internal networks do not run RIP. They use static routing because they are small networks. This means the FortiGate devices have to redistribute any static routes they learn so that the internal networks can communicate with each other.


Where possible in this example, the default values will be used (or the most general settings). This is intended to provide an easier configuration that will require less troubleshooting.

In this example, the routers, networks, interfaces used, and IP addresses are as follows. Note that the interfaces that connect Router2 and Router3 also connect to the R&D network.

RIP example network topology

Network	Router	Interface & alias	IP address		
Sales	Router1	port1 (internal)	10.11.101.101		
		port2 (router2)	10.11.201.101		
		port3 (router3)	10.11.202.101		
		port1 (internal)	10.12.101.102		
	Router2	port2 (router1)	10.11.201.102		
		port3 (router4)	10.14.201.102		
		port4 (ISP)	172.20.120.102		
		R&D	Router3	port1 (internal)	10.12.101.103
port2 (router1)	10.11.201.103				
port3 (router4)	10.14.202.103				
port4 (ISP)	172.20.120.103				
Router4	port1 (internal)		10.14.101.104		
	Accounting		Router4	port2 (router2)	10.14.201.104
	port3 (router3)		10.14.202.104		

Network topology for the simple RIP example



Assumptions

This example makes the following assumptions:

- All FortiGate devices have 5.0 firmware and are running factory default settings.
- All CLI and GUI navigation assumes the unit is running in NAT/Route operating mode, with VDOMs disabled.
- All FortiGate devices have interfaces labeled port1 through port4, as required.
- All firewalls have been configured for each FortiGate to allow the required traffic to flow across interfaces.
- Only FortiGate devices are running RIP on the internal networks.
- Router2 and Router3 are connected through the internal network for R&D.
- Router2 and Router3 each have their own connection to the Internet, indicated in black in the diagram above.

General configuration steps

This example is very straightforward. The steps involved are:

- [Configuring FortiGate system information](#)
- [Configuring FortiGate RIP router information](#)
- [Configuring other networking devices](#)
- [Testing network configuration](#)

Configuring FortiGate system information

You must configure the hostname and interfaces for each FortiGate.

For IP numbering, Router2 and Router3 use the numbering for the other routers, where needed.

Router2 and Router3 have link health monitoring enabled on the ISP interfaces using Ping. Remember to contact the ISP and confirm their server has ping enabled.

Configure the hostname, interfaces, and default route

To configure Router1 system information - GUI:

1. Go to **System > Settings**.
2. In the **Host name** field, enter `Router1`.
3. Go to **Network > Static Routes**.
4. Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.5/255.255.255.0
Interface	port2 (router2)
Administrative Distance	40

5. Enter a second default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.5/255.255.255.0
Interface	port3 (router3)
Administrative Distance	40

6. Go to **Network > Interfaces**.
7. Edit port1 (internal) interface.
8. Set the following information, and select **OK**.

Alias	internal
IP/Network Mask	10.11.101.101/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Internal sales network
Interface State	Enabled

9. Edit port2 (router2) interface.
10. Set the following information, and select **OK**.

Alias	router2
IP/Network Mask	10.11.201.101/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Link to R&D network & Internet through Router2
Interface State	Enabled

11. Edit port3 (router3) interface.

12. Set the following information, and select **OK**.

Alias	router3
IP/Network Mask	10.11.202.101/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Link to R&D network and Internet through Router3
Interface State	Enabled

To configure Router1 system information - CLI:

```

config system global
    set hostname Router1
end

config router static
    edit 1
        set device "port2"
        set distance 45
        set gateway 10.11.201.102
    next
    edit 2
        set device "port3"
        set distance 45
        set gateway 10.11.202.103
    end
end

config system interface
    edit port1
        set alias internal
        set ip 10.11.101.101/255.255.255.0
        set allowaccess https ssh ping
        set description "Internal sales network"
    next
    edit port2
        set alias ISP
        set allowaccess https ssh ping
        set ip 10.11.201.101/255.255.255.0

```



```

        set description "Link to R&D network & Internet through Router2"
    next
    edit port3
        set alias router3
        set ip 10.11.202.101/255.255.255.0
        set allowaccess https ssh ping
        set description "Link to R&D network & Internet through Router2"
    end
end
end

```

To configure Router2 system information - GUI:

1. Go to **System > Settings**.
2. In the **Host name** field, enter `Router2`.
3. Go to **Network > Static Routes**.
4. Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.5/255.255.255.0
Interface	port4 (ISP)
Administrative Distance	5

5. Go to **Network > Interfaces**.
6. Edit port1 (internal) interface.
7. Set the following information and select **OK**.

Alias	internal
IP/Network Mask	10.12.101.102/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	R&D internal network and Router3
Interface State	Enabled

8. Edit port2 (router1) interface.
9. Set the following information and select **OK**.

Alias	router1
IP/Network Mask	10.12.201.102/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Link to Router1 and the Sales network
Interface State	Enabled

10. Edit port3 (router4) interface.
11. Set the following information and select **OK**.

Alias	router4
IP/Network Mask	10.12.301.102/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Link to Router4 and the accounting network
Interface State	Enabled

12. Edit port4 (ISP) interface.
13. Set the following information and select **OK**.

Alias	ISP
IP/Network Mask	172.20.120.102/255.255.255.0
Administrative Access	HTTPS SSH PING
Detect and Identify Devices	enable
Comments	Internet through ISP
Interface State	Enabled

To configure Router2 system information - CLI:

```

config system global
    set hostname Router2
end
config router static
    edit 1
        set device "port4"
        set distance 5
        set gateway 172.20.130.5
    end
end
config system interface
    edit port1
        set alias internal
        set ip 10.11.101.102/255.255.255.0
        set allowaccess https ssh ping
        set description "Internal RnD network and Router3"
    next
    edit port2
        set alias router1
        set allowaccess https ssh ping
        set ip 10.11.201.102/255.255.255.0
        set description "Link to Router1"
    next

```

```

edit port3
    set alias router3
    set ip 10.14.202.102/255.255.255.0
    set allowaccess https ssh ping
    set description "Link to Router4"
next
edit port4
    set alias ISP
    set ip 172.20.120.102/255.255.255.0
    set allowaccess https ssh ping
    set description "ISP and Internet"
end
end

```

To configure Router3 system information - GUI:

1. Go to **System > Settings**.
2. In the **Host name** field, enter Router3.
3. Go to **Network > Static Routes**.
4. Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.5/255.255.255.0
Interface	port4 (ISP)
Administrative Distance	5

5. Go to **Network > Interfaces**.
6. Edit port1 (internal) interface.
7. Set the following information and select **OK**.

Alias	internal
IP/Network Mask	10.12.101.103/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	R&D internal network and Router2
Interface State	Enabled

8. Edit port2 (router1) interface.
9. Set the following information and select **OK**.

Alias	router1
IP/Network Mask	10.13.201.103/255.255.255.0

Administrative Access	HTTPS SSH PING
Comments	Link to Router1 and Sales network
Interface State	Enabled

10. Edit port3 (router4) interface.

11. Set the following information and select **OK**.

Alias	router4
IP/Network Mask	10.13.301.103/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Link to Router4 and accounting network
Interface State	Enabled

12. Edit port4 (ISP) interface.

13. Set the following information and select **OK**.

Alias	ISP
IP/Network Mask	172.20.120.103/255.255.255.0
Administrative Access	HTTPS SSH PING
Detect and Identify Devices	enable
Comments	Internet and ISP
Interface State	Enabled

To configure Router3 system information - CLI:

```

config system global
    set hostname Router3
end
config router static
    edit 1
        set device "port4"
        set distance 5
        set gateway 172.20.130.5
    end
end
config system interface
    edit port1
        set alias internal
        set ip 10.12.101.103/255.255.255.0
        set allowaccess https ssh ping
        set description "Internal RnD network and Router2"
    end
end

```

```

next
edit port2
    set alias ISP
    set allowaccess https ssh ping
    set ip 10.11.201.103/255.255.255.0
    set description "Link to Router1"
next
edit port3
    set alias router3
    set ip 10.14.202.103/255.255.255.0
    set allowaccess https ssh ping
    set description "Link to Router4"
next
edit port4
    set alias ISP
    set ip 172.20.120.103/255.255.255.0
    set allowaccess https ssh ping
    set description "ISP and Internet"
end
end

```

To configure Router4 system information - GUI:

1. Go to **System > Settings**.
2. In the **Host name** field, enter Router4.
3. Go to **Network > Static Routes**.
4. Edit the default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.5/255.255.255.0
Interface	port2 (router2)
Administrative Distance	40

5. Enter a second default route and enter the following information:

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.20.120.5/255.255.255.0
Interface	port3 (router3)
Administrative Distance	40

6. Go to **Network > Interfaces**.
7. Edit port 1 (internal) interface.
8. Set the following information and select **OK**.

Alias	internal
--------------	----------

IP/Network Mask	10.14.101.104/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Internal accounting network
Interface State	Enabled

9. Edit port 2 (router2) interface.
10. Set the following information and select **OK**.

Alias	router2
IP/Network Mask	10.14.201.104/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Link to R&D network & Internet through Router2
Interface State	Enabled

11. Edit port 3 (router3) interface.
12. Set the following information and select **OK**.

Alias	router3
IP/Network Mask	10.14.301.104/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Link to R&D network and Internet through Router3
Interface State	Enabled

To configure Router4 system information - CLI:

```

config system global
    set hostname Router4
end
config router static
    edit 1
        set device "port2"
        set distance 45
        set gateway 10.14.201.102
    next
    edit 2
        set device "port3"
        set distance 45
        set gateway 10.14.202.103
    end
end
config system interface

```

```
edit port1
    set alias internal
    set ip 10.14.101.104/255.255.255.0
    set allowaccess https ssh ping
    set description "Internal sales network"
next
edit port2
    set alias router2
    set allowaccess https ssh ping
    set ip 10.14.201.104/255.255.255.0
    set description "Link to R&D network & Internet through Router2"
next
edit port3
    set alias router3
    set ip 10.14.202.104/255.255.255.0
    set allowaccess https ssh ping
    set description "Link to R&D network & Internet through Router2"
end
end
```

Configuring FortiGate RIP router information

With the interfaces configured, RIP can now be configured on the FortiGate.

For each FortiGate, the following steps will be taken:

- Configure RIP version used
- Redistribute static networks
- Add networks serviced by RIP
- Add interfaces that support RIP on the FortiGate

Router1 and Router4 are configured the same. Router2 and Router3 are configured the same. These routers will be grouped accordingly for the following procedures. Repeat the procedures once for each FortiGate.

Configure RIP settings on Router1 and Router4 - GUI:

1. Go to **Network > RIP**.
2. Select **2** for **Version**.
3. In **Advanced Options**, under **Redistribute** enable **Static**. Leave the other advanced options at their default values.
4. Under **Networks**, add the following networks:
 - 10.11.0.0/255.255.0.0
 - 10.12.0.0/255.255.0.0
 - 10.14.0.0/255.255.0.0
 - 172.20.120.0/255.255.255.0
6. Under **Interfaces**, select **Create New** and set the following information:

Interface	port1 (internal)
Passive	disabled

Authentication	None
Send Version	Both
Receive Version	Both

7. Under **Interfaces** select **Create New** and set the following information:

Interface	port2 (router2)
Passive	disabled
Authentication	None
Send Version	Both
Receive Version	Both

8. Under **Interfaces**, select **Create New** and set the following information:

Interface	port3 (router3)
Passive	disabled
Authentication	None
Send Version	Both
Receive Version	Both

Configure RIP settings on Router1 and Router4 - CLI:

```

config router rip
  set version 2
  config interface
    edit "port1"
      set receive-version 1 2
      set send-version 1 2
    next
    edit "port2"
      set receive-version 1 2
      set send-version 1 2
    next
    edit "port3"
      set receive-version 1 2
      set send-version 1 2
    end
  config network
    edit 1
      set prefix 10.11.0.0 255.255.0.0
    next
    edit 2
      set prefix 10.12.0.0 255.255.0.0

```



```

    next
    edit 3
        set prefix 10.14.0.0 255.255.0.0
    next
    edit 4
        set prefix 172.20.120.0 255.255.255.0
    end
    config redistribute "static"
        set status enable
    end
end
end

```

Configure RIP settings on Router2 and Router3 - GUI:

1. Go to **Network > RIP**.
2. Select **2** for **RIP**.
3. In **Advanced Options**, under **Redistribute** enable **Static**. Leave the other advanced options at their default values.
4. Under **Networks**, add the following networks:
 - 10.11.0.0/255.255.0.0
 - 10.12.0.0/255.255.0.0
 - 10.14.0.0/255.255.0.0
 - 172.20.120.0/255.255.255.0
6. Under **Interfaces**, select **Create New** and set the following information:

Interface	port1 (internal)
Passive	disabled
Authentication	None
Send Version	Both
Receive Version	Both

7. Under **Interfaces**, select **Create New** and set the following information:

Interface	port2 (router1)
Passive	disabled
Authentication	None
Send Version	Both
Receive Version	Both

8. Under **Interfaces**, select **Create New** and set the following information:

Interface	port3 (router4)
Passive	disabled
Authentication	None
Send Version	Both
Receive Version	Both

9. Under **Interfaces**, select **Create New** and set the following information:

Interface	port4 (ISP)
Passive	disabled
Authentication	None
Send Version	Both
Receive Version	Both

Configure RIP settings on Router2 and Router3 - GUI:

```

config router rip
  set version 2
  config interface
    edit "port1"
      set receive-version 1 2
      set send-version 1 2
    next
    edit "port2"
      set receive-version 1 2
      set send-version 1 2
    next
    edit "port3"
      set receive-version 1 2
      set send-version 1 2
    end
    edit "port4"
      set receive-version 1 2
      set send-version 1 2
    end
  config network
    edit 1
      set prefix 10.11.0.0 255.255.0.0
    next
    edit 2
      set prefix 10.12.0.0 255.255.0.0
    next
    edit 3
      set prefix 10.14.0.0 255.255.0.0
    next
    edit 4

```

```
        set prefix 172.20.120.0 255.255.255.0
    end
    config redistribute "static"
        set status enable
    end
end
```

Configuring other networking devices

In this example, there are two groups of other devices on the the network: internal devices and the ISP.

The first is the internal network devices on the Sales, R&D, and Accounting networks. This includes simple static routers, computers, printers, and other network devices. Once the FortiGate devices are configured, the internal static routers need to be configured using the internal network IP addresses. Otherwise, there should be no configuration required.

The second group of devices is the ISP. This consists of the RIP router the FortiGate Router2 and Router3 connect to. You need to contact your ISP and ensure they have your information for your network, such as the IP addresses of the connecting RIP routers, what version of RIP your network supports, and what authentication (if any) is used.

Testing network configuration

Once the network has been configured, you need to test that it works as expected.

The two series of tests you need to run are to test the internal networks can communicate with each other, and that the internal networks can reach the Internet.

Use ping, traceroute, and other networking tools to run these tests.

If you encounter problems, for troubleshooting help consult ["Troubleshooting RIP" on page 2121](#).

IPsec auto discovery support

The following routing settings are available in the CLI to support IPsec auto discovery. They are designed for:

- Supporting the RIPvng (RIP next generation) network command
- Limiting the maximum metric allowed to output for RIPvng
- Fix NSM missing kernel address update information

The actual new settings are:

```
config router rip
    set max-out-metric <integer value 1 - 15>
end

config router ripng
    set max-out-metric <integer value 1 - 15>
end

config router ripng
    config network
        edit <network-ID>
            set prefix <IPv6-prefix>
        end
    end
```

RIPng: RIP and IPv6

RIP next generation, or RIPng, is the version of RIP that supports IPv6.

This is an example of a typical small network configuration using RIPng routing.

Your internal R&D network is working on a project for a large international telecom company that uses IPv6. For this reason, you have to run IPv6 on your internal network and you have decided to use only IPv6 addresses.

Your network has two FortiGate devices running the RIPng dynamic routing protocol. Both FortiGate devices are connected to the ISP router and the internal network. This configuration provides some redundancy for the R&D internal network, allowing it to reach the Internet at all times.

Network layout and assumptions

Basic network layout

Your internal R&D network is working on a project for a large international telecom company that uses IPv6. For this reason, you have to run IPv6 on your internal network and you have decided to use only IPv6 addresses.

Your network has two FortiGate devices running the RIPng dynamic routing protocol. Both FortiGate devices are connected to the ISP router and the internal network. This configuration provides some redundancy for the R&D internal network, allowing it to reach the Internet at all times.

All internal computers use RIP routing, so no static routing is required. And all internal computers use IPv6 addresses.


Where possible in this example, the default values will be used (or the most general settings). This is intended to provide an easier configuration that will require less troubleshooting.

In this example, the routers, networks, interfaces used, and IP addresses are as follows.

Example RIP network topology

Network	Router	Interface & alias	IPv6 address
R&D	Router1	port1 (internal)	2002:A0B:6565:0:0:0:0:0
		port2 (ISP)	2002:AC14:7865:0:0:0:0:0
	Router2	port1 (internal)	2002:A0B:6566:0:0:0:0:0
		port2 (ISP)	2002:AC14:7866:0:0:0:0:0

Example network topology for IPV6 RIPng



Assumptions

The following assumptions have been made concerning this example.

- All FortiGate devices have 5.0 firmware and are running factory default settings.
- All CLI and GUI navigation assumes the unit is running in NAT/Route operating mode, with VDOMs disabled.
- All FortiGate devices have interfaces labeled port1 and port2, as required.
- All firewalls have been configured for each FortiGate to allow the required traffic to flow across interfaces.
- All network devices support IPv6 and are running RIPng.

Configuring the FortiGate system information

Each FortiGate needs IPv6 enabled, a new hostname, and interfaces configured.

To configure system information on Router1 - GUI:

1. Go to **System > Dashboard > Status**.
2. For **Host name**, select **Change**.
3. Enter "Router1".
4. Go to **System > Config > Features**.

5. In **Basic Features**, enable **IPv6**, and select **Apply**.
6. Go to **System > Network > Interfaces**.
7. Edit port1 (internal) interface.
8. Set the following information and select **OK**.

Alias	internal
IP/Network Mask	2002:A0B:6565::/0
Administrative Access	HTTPS SSH PING
Description	Internal RnD network
Administrative Status	Up

9. Edit port2 (ISP) interface.
10. Set the following information and select **OK**.

Alias	ISP
IP/Network Mask	2002:AC14:7865::/0
Administrative Access	HTTPS SSH PING
Description	ISP and Internet
Administrative Status	Up

To configure system information on Router1 - CLI:

```

config system global
    set hostname Router1
    set gui-ipv6 enable
end
config system interface
    edit port1
        set alias internal
        set allowaccess https ping ssh
        set description "Internal RnD network"
        config ipv6
            set ip6-address 2002:a0b:6565::/0
        end
    next
    edit port2
        set alias ISP
        set allowaccess https ping ssh
        set description "ISP and Internet"
        config ipv6
            set ip6-address 2002:AC14:7865::
        end
    end
end

```

To configure system information on Router2 - GUI:

1. Go to **System > Dashboard > Status**.
2. For **Host name**, select **Change**.
3. Enter "Router2".
4. Go to **System > Config > Features**.
5. In **Basic Features**, enable **IPv6**, and select **Apply**.
6. Go to **System > Network > Interfaces**.
7. Edit port1 (internal) interface.
8. Set the following information and select **OK**.

Alias	internal
IP/Network Mask	2002:A0B:6566::/0
Administrative Access	HTTPS SSH PING
Description	Internal RnD network
Administrative Status	Up

9. Edit port2 (ISP) interface.
10. Set the following information and select **OK**.

Alias	ISP
IP/Network Mask	2002:AC14:7866::/0
Administrative Access	HTTPS SSH PING
Description	ISP and Internet
Administrative Status	Up

To configure system information on Router2 - CLI:

```

config system global
    set hostname Router2
    set gui-ipv6 enable
end
config system interface
    edit port1
        set alias internal
        set allowaccess https ping ssh
        set description "Internal RnD network"
        config ipv6
            set ip6-address 2002:a0b:6566::/0
        end
    next
    edit port2
        set alias ISP
        set allowaccess https ping ssh

```

```
set description "ISP and Internet"
config ipv6
    set ip6-address 2002:AC14:7866::
end
end
```

Configuring RIPng on FortiGate

Now that the interfaces are configured, you can configure RIPng on the FortiGate devices.

There are only two networks and two interfaces to include: the internal network and the ISP network. There is no redistribution and no authentication. In RIPng there is no specific command to include a subnet in the RIP broadcasts. There is also no information required for the interfaces beyond including their name.

As this is a CLI only configuration, configure the ISP router and the other FortiGate as neighbors. This was not part of the previous example as this feature is not offered in the GUI. Declaring neighbors in the configuration like this will reduce the discovery traffic when the routers start up.

Since RIPng is not supported in the GUI, this section will only be entered in the CLI.

To configure RIPng on Router1 - CLI:

```
config router ripng
config interface
    edit port1
    next
    edit port2
end
config neighbor
    edit 1
        set interface port1
        set ipv6 2002:a0b:6566::/0
    next
    edit 2
        set interface port2
        set ipv6 2002:AC14:7805::/0
end
```

To configure RIPng on Router2 - CLI:

```
config router ripng
config interface
    edit port1
    next
    edit port2
end
config neighbor
    edit 1
        set interface port1
        set ipv6 2002:a0b:6565::/0
    next
    edit 2
        set interface port2
        set ipv6 2002:AC14:7805::/0
end
```


Configuring other network devices

The other devices on the internal network all support IPv6 and are running RIPng, where applicable. They only need to know the internal interface network addresses of the FortiGate devices.

The ISP routers need to know the FortiGate information, such as IPv6 addresses.

Testing the configuration

In addition to normal testing of your network configuration, you must also test the IPv6 part of this example.

For troubleshooting problems with your network, see the *Troubleshooting Handbook*.

For troubleshooting problems with RIP, see "[Troubleshooting RIP](#)" on page 2121.

Testing the IPv6 RIPng information

There are some commands to use when checking that your RIPng information is correct on your network. These are useful to check on your RIPng FortiGate devices on your network. Comparing the output between devices will help you understand your network better, and also track down any problems:

```
diagnose ipv6 address list
```

View the local scope IPv6 addresses used as next-hops by RIPng on the FortiGate:

```
diagnose ipv6 route list
```

View ipv6 addresses that are installed in the routing table:

```
get router info6 routing-table
```

View the routing table. This information is almost the same as the previous `diagnose ipv6 route list` command, but it is presented in a format that is easier to read.

```
get router info6 rip interface external
```

View the brief output on the RIP information for the interface listed. This includes information such as, if the interface is up or down, what routing protocol is being used, and whether passive interface or split horizon are enabled.

```
get router info6 neighbor-cache list
```

View the IPv6/MAC address mapping. This also displays the interface index and name associated with the address.

OSPF

Open Shortest Path First (OSPF) is a link-state interior routing protocol that is widely used in large enterprise organizations. It only routes packets within a single autonomous system (AS). This is different from BGP, because BGP can communicate between ASs.

OSPF version 2 (OSPFv2) was defined in 1998 in [RFC 2328](#). OSPF was designed to support classless IP addressing and variable subnet masks. This was a shortcoming of the earlier RIP protocols.

Updates to OSPFv2 are included in OSPF version 3 (OSPFv3), defined in 2008 in [RFC 5340](#). OSPFv3 includes support for IPv6 addressing, where OSPF2 only supports IPv4 addressing.

The main benefit of OSPF is that it detects link failures in the network quickly and within seconds, has converged network traffic successfully without any networking loops. Also, OSPF has many features to control which routes are propagated and which are not, maintaining smaller routing tables. OSPF can also provide better load-balancing on external links than other interior routing protocols.

The parts and terminology of OSPF

The parts and terminology of OSPF include the following sections.

OSPFv3 and IPv6

OSPF version 3 (OSPFv3) includes support for IPv6. Generally, all IP addresses are in IPv6 format instead of IPv4. However, OSPFv3 area numbers use the same 32-bit numbering system as OSPFv2, as described in [RFC 2740](#). Likewise, the router ID and area ID are in the same format as OSPFv2.

As with most advanced routing features on a FortiGate, IPv6 settings for dynamic routing protocols must be enabled before they are visible in the GUI. To enable IPv6 configuration in the GUI, enable it in **System > Feature Visibility**.

For IPv6, the main difference in OSPFv3 is that rather than using a network statement to enable OSPFv3 on an interface, you define OSPF6 (OSPF for IPv6) interfaces, which are bound to the interface and area. This configuration must be done in the CLI, as follows (with sample interfaces and addresses):

```
config router ospf6
  config area
    edit 0.0.0.0
  next
end
config ospf6-interface
  edit "tunnel"
    set interface "to_FGT300A-7"
  next
  edit "internal_lan"
    set interface "port1"
  next
  set router-id 10.174.0.113
end
```

Note that OSPFv3 neighbors use link-local IPv6 addresses, but with broadcast and point-to-point network types, and neighbors are automatically discovered. You only have to manually configure neighbors when using non-broadcast network types.

Router ID

In OSPF, each router has a unique 32-bit number that is called its router ID. Often, this 32-bit number is written the same as a 32-bit IPv4 address would be written in dotted decimal notation. However, some brands of routers, such as Cisco routers, support a router ID entered as an integer instead of an IP address.

It's a good idea not to use an IP address for the router ID that is already in use on the router. The router ID doesn't have to be a particular IP address on the router. By choosing a different number, it's harder to get confused about which number you're looking at. It's a good idea to use as many of the area's numbers as possible. For example, if you have 15 routers in area 0.0.0.0, they could be numbered from 0.0.0.1 to 0.0.0.15. If you have an area 1.1.1.1, then routers in that area could start at 1.1.1.10.

You can manually set the router ID on a FortiGate:

To manually set an OSPF router ID of 0.0.1.1 - GUI:

1. Go to **Router > Dynamic > OSPF**.
2. For **Router ID**, enter 0.0.1.1.
3. Select **Apply**.

To manually set an OSPF router ID of 0.0.1.1 - CLI:

```
config router ospf
  set router-id 0.0.1.1
end
```

Adjacency

In an OSPF routing network, an OSPF router sends out OSPF hello packets when it boots up, to try to find any neighbors (routers that have access to the same network as the router booting up). Once neighbors are discovered and Hello packets are exchanged, updates are sent and the link state databases of both neighbors are synchronized. At this point, these neighbors are said to be adjacent.

For two OSPF routers to become neighbors, the following conditions must be met:

- The subnet mask used on both routers must be the same subnet.
- The subnet number derived using the subnet mask and each router's interface IP address must match.
- The hello interval and the dead interval must match.
- The routers must have the same OSPF area ID. If they're in different areas, they're not neighbors.
- If authentication is used, they must pass authentication checks.

If any of these parameters are different between the two routers, the routers do not become OSPF neighbors and can't be adjacent. If the routers become neighbors, they're adjacent.

Adjacency and neighbors

Neighbor routers can be in a two-way state, and not be adjacent. Adjacent routers normally have a neighbor state of FULL. Neighbors only exchange hello packets and don't exchange routing updates. Adjacent routers exchange LSAs (LSDB information) as well as hello packets. A good example of an adjacent pair of routers is the designated router (DR) and backup designated router (BDR).

You can check on the state of an OSPF neighbor using the CLI `get router info ospf neighbor all` command. For more information, see ["Checking the state of OSPF neighbors" on page 2159](#).

Why adjacency is important

It's important to have adjacent pairs of routers in the OSPF routing domain because routing protocol packets are only passed between adjacent routers. This means adjacency is required for two OSPF routers to exchange routes.

If there's no adjacency between two routers, such as one on the 172.20.120.0 network and another on the 10.11.101.0 network, the routers don't exchange routes. This makes sense because if all OSPF routers on the OSPF domain exchanged updates, it would flood the network.

Also, it's better for updates to progress through adjacent routers to ensure there are no outages along the way. Otherwise, updates could skip over routers that are potentially offline, causing longer routing outages and delays, while the OSPF domain learns of this outage later on.

If the OSPF network has multiple border routers and multiple connections to external networks, the designated router (DR) determines which router pairs become adjacent. The DR can accomplish this because it maintains the complete topology of the OSPF domain, including which router pairs are adjacent.

The backup designated router (BDR) also has this information in case the DR goes offline.

Designated router and backup router

In OSPF, a router can have a number of different roles to play.

A designated router (DR) is the designated broadcasting router interface for an AS. It looks after all of the initial contact and other routing administration traffic. Having only one router do all of this greatly reduces the network traffic and collisions.

If something happens and the designated router goes offline, the backup designated router (BDR) takes over. An OSPF interface on a FortiGate can become either a DR or BDR. Both the DR and the BDR cover the same area and are elected at the same time. The election process doesn't have many rules, but the exceptions can become complex.

Benefits

The OSPF concept of the designated router is a big step above RIP. With all RIP routers doing their own updates all the time, RIP suffers from frequent and sometimes unnecessary updates that can slow down your network. With OSPF, not only do routing changes only happen when a link state changes instead of any tiny change to the routing table, but the designated router reduces this overhead traffic even more.

However, smaller network topologies may have only a couple of routers besides the designated router. This may seem excessive, but it maintains the proper OSPF form and it still reduces the administration traffic, but to a lesser extent than on a large network. Also, your network topology is ready whenever you choose to expand your network.

DR and BDR election

An election chooses DR and BDR from all the available routers. The election is primarily based on the priority setting of the routers, where the highest priority becomes the DR and the second highest becomes the BDR. To resolve any ties, the router with the highest router ID wins. For example, a router with a router ID of 192.168.0.1 would win over a router with a router ID of 10.1.1.2.

The router priority can vary from 0 to 255, but at 0 a router won't become a DR or BDR. If a router with a higher priority comes online after the election, it must wait until the DR and BDR go offline before it becomes the DR.

If the original DR goes offline, but is then available when the BDR goes offline later on, the original DR will be promoted back to DR without an election leaving the new BDR as it is.

With the FortiGate, to configure the port1 interface to be a potential OSPF DR or BDR called `ospf_DR` on the network, you need to raise the priority of the router to a very high number, such as 250 out of 255. This ensures the interface has a chance to be a DR, but won't guarantee that it'll be one. To help ensure it becomes a DR, you should give the interface a low numbered IP address, such as 10.1.1.1 instead of 192.168.1.1 (but that isn't part of this example). Enter the following command:

```
config router ospf
  config ospf-interface
    edit "ospf_DR"
      set priority 250
    next
  end
```

Area

An OSPF area is a smaller part of the larger OSPF AS. Areas are used to limit the link state updates that are sent out. The flooding used for these updates would overwhelm a large network, so it's divided into these smaller areas for manageability.

If there are two or more routers that are viable within an area, there will always be a designated router (DR) and a backup designated router (BDR). For more information about these router roles, see ["Designated router and backup router" on page 2149](#).

Defining a private OSPF area involves the following:

- Assigning a 32-bit number to the area that is unique on your network
- Defining the characteristics of one or more OSPF areas
- Creating associations between the OSPF areas that you defined and the local networks to include in the OSPF area
- Adjusting the settings of OSPF-enabled interfaces, if required



IPv6 OSPF area numbers use the same 32-bit number notation as IPv4 OSPF.

If you are using the GUI to perform these tasks, follow the procedures summarized below.

FortiGate devices support the four main types of OSPF areas: backbone, stub, not-so-stubby, and regular.

Backbone area

Every OSPF network has at least one AS and every OSPF network has a backbone area. The backbone is the main area, and possibly the only area. All other OSPF areas are connected to a backbone area. This means if two areas want to pass routing information back and forth, that routing information will go through the backbone on its way between those areas. For this reason, the backbone not only has to connect to all other areas in the network, but also has to be uninterrupted in order to be able to pass traffic to all points of the network.

The backbone area is referred to as area 0 because it has an IP address of 0.0.0.0.

Stub area

A stub area is an OSPF area that receives no outside routes advertised into it. All routing in it is based on a default route. This essentially isolates it from outside areas.

Stub areas are useful for small networks that are part of a larger organization, especially if the networking equipment can't handle routing large amounts of traffic passing through, or if there are other reasons to prevent outside traffic, such as security. For example, most organizations don't want their accounting department to be the center of their network with everyone's traffic passing through there. It increases the security risks, slows down the network, and it generally doesn't make sense.

A variation on the stub area is the totally stubby area. It's a stub area that doesn't allow summarized routes.

NSSA

A not-so-stubby-area (NSSA) is a stub area that allows for external routes to be injected into it. While it still doesn't allow routes from external areas, it's not limited to using only the default route for internal routing.

Regular area

A regular area is what all the other ASs are, all the non-backbone, non-stub, and non-NSSA areas. A regular area generally has a connection to the backbone, does receive advertisements of outside routes, and doesn't have an area number of 0.0.0.0.

Authentication

In the OSPF packet header, there are two authentication-related fields: AuType and Authentication.

All OSPF packet traffic is authenticated. Multiple types of authentication are supported in OSPFv2. However, in OSPFv3, there's no authentication built-in but it's assumed that IPsec is used for authentication instead.

Packets that fail authentication are discarded.

Null authentication

Null authentication indicates there's no authentication being used. In this case, the 16-byte authentication field isn't checked, and can be any value. However, checksumming is still used to locate errors. On a FortiGate, this is the `none` option for authentication.

Simple password authentication

Simple password refers to a standard plain text string of characters. The same password is used for all transactions on a network. The main use for this type of authentication is to prevent routers from accidentally joining the network. Simple password authentication is vulnerable to many forms of attack, and isn't recommended as a secure form of authentication.

Cryptographic authentication

Cryptographic authentication involves the use of a shared secret key to authenticate all router traffic on a network. The key is never sent over the network in the clear. A packet is sent and a condensed and encrypted form of the packet is appended to the end of the packet. A non-repeating sequence number is included in the OSPF packet to protect against replay attacks that could try to use already sent packets to disrupt the network. When a packet is accepted as authentic, the authentication sequence number is set to the packet sequence number. If a replay attack is attempted, the packet sent will be out of sequence and ignored.

A FortiGate supports all three levels of authentication through the authentication keyword associated with creating an OSPF interface .

For example, to create an OSPF interface called `Accounting` on the port1 interface that is a broadcast interface, has a hello interval of 10 seconds, has a dead interval of 40 seconds, uses text authentication (simple password) with a password of "ospf_test", enter the following CLI commands:

```
config router ospf
  config ospf-interface
    edit Accounting
      set interface port1
      set network-type broadcast
      set hello-interval 10
      set dead-interval 40
      set authentication text
      set authentication-key ospf_test
    next
  end
```

Hello and dead intervals

The OSPF Hello protocol is used to discover and maintain communications with neighboring routers.

Hello packets are sent out at a regular interval for this purpose. The DR sends out the hello packets. In a broadcast network, the multicast address of 224.0.0.5 is used to send out hello packets. New routers on the network listen for and reply to these packets to join the OSPF area. If a new router never receives a hello packet, other routers won't know it is there and won't communicate with it. However, once a new router is discovered, the DR adds it to the list of routers in that area and it's integrated into the routing calculations.

Dead interval is the time other routers wait before declaring a neighbor dead (offline). It's very important to set a reasonable dead interval. If this interval is too short, routers will be declared offline when they are just slow or momentarily inaccessible, and link state updates will happen more than they need to, using more bandwidth. If the dead interval is too long, it will slow down network traffic overall if online routers attempt to contact offline ones instead of re-routing traffic.

FortiOS also supports OSPF fast-hello, which provides a way to send multiple hello packets per second. This is achieved by setting a dead-interval to one second. The hello-multiplier, which can be any number between 4 and 10, determines the number of hello packets that will be sent every second. The CLI syntax for OSPF fast-hello is the following:

```
config ospf-interface
edit ospf1
set interface port1
set network-type broadcast
set dead-interval 1
set hello-multiplier 4
next
end
```

Access lists

Access lists are filters used by OSPF routing on a FortiGate. An access list provides a list of IP addresses and the action to take for them. An access list essentially makes it easy to group addresses that will be treated the same into the same group, independent of their subnets or other matching qualities. You add a rule for each address or subnet that you want to include, specifying the action to take for it. For example, if you want all traffic from one department to be routed a particular way, even in different buildings, you can add all the addresses to an access list and then handle that list all at once.

Each rule in an access list consists of a prefix (IP address and netmask), the action to take for this prefix (permit or deny), and whether to match the prefix exactly or to match the prefix and any more specific prefix.

The FortiGate attempts to match a packet against the rules in an access list, starting at the top of the list. If it finds a match for the prefix, it takes the action specified for that prefix. If no match is found, the default action is deny.

Access lists greatly speed up configuration and network management. When there is a problem, you can check each list instead of individual addresses. It also eases troubleshooting because if all addresses on one list have problems, it eliminates many possible causes right away.

If you are using the OSPF+ IPv6 protocols, you will need to use access-list6, the IPv6 version of access list. The only difference is that access-list6 uses IPv6 addresses.

For example, if you want to create an access list called `test_list` that only allows an exact match of 10.10.10.10 and 11.11.11.11, enter the following CLI commands:

```
config router access-list
```

```

edit test_list
config rule
edit 1
set prefix 10.10.10.10 255.255.255.255
set action allow
set exact-match enable
next
edit 2
set prefix 11.11.11.11 255.255.255.255
set action allow
set exact-match enable
next
next
end

```

Another example is if you want to deny ranges of addresses in IPv6 that start with the IPv6 equivalents of 10.10.10.10 and 11.11.11.11, enter the following `access-list6` CLI commands:

```

config router access-list6
edit test_list_ip6
config rule
edit 1
set prefix6 2002:A0A:A0A:0:0:0:0:0/48
set action deny
next
edit 2
set prefix6 2002:B0B:B0B:0:0:0:0:0/48
set action deny
next
next
end

```

To use an `access_list`, you must call it from a routing protocol such as RIP. The following example uses the `access_list` from the earlier example called `test_list` to match routes coming in on the `port1` interface. When there's a match, it'll add 3 to the hop count metric for those routes to artificially decrease their priority. Enter the following CLI commands:

```

config router ospf
config distribute-list
edit 5
set access-list test_list
set protocol connected
next
end

```

If you're setting a prefix of 128.0.0.0, use the format 128.0.0.0/1. The default route 0.0.0.0/0 can't be matched exactly with an `access-list`. A `prefix-list` must be used for this purpose.

How OSPF works

An OSPF installation consists of one or more areas. An OSPF area is typically divided into logical areas linked by Area Border Routers (ABR). A group of contiguous networks form an area. An ABR links one or more areas to the OSPF network backbone (area ID 0). For more information, see ["Dynamic routing" on page 2093](#).

OSPF is an interior routing protocol. It includes a backbone AS and possibly additional ASs. The DR and BDR are elected from potential routers with the highest priorities. The DR handles much of the administration to lower the

network traffic required. New routers are discovered through hello packets sent from the DR using the multicast address of 224.0.0.5. If the DR goes offline at any time, the BDR has a complete table of routes that it uses when it takes over as the DR router.

OSPF doesn't use UDP or TCP, but is encapsulated directly in IP datagrams as protocol 89. This is in contrast to RIP and BGP. OSPF handles its own error detection and correction functions.

The OSPF protocol, when running on IPv4, can operate securely between routers, optionally using a variety of authentication methods to allow only trusted routers to participate in routing. OSPFv3, running on IPv6, no longer supports protocol-internal authentication. Instead, it relies on IPv6 protocol security (IPsec).

Other important parts of how OSPF works include:

- [OSPF router discovery](#)
- [How OSPF works on FortiGate devices](#)
- [External routes](#)
- [Link state database and route updates](#)
- [OSPF packets](#)

OSPF router discovery

OSPF-enabled routers generate link state advertisements (LSA) and send them to their neighbors whenever the status of a neighbor changes or a new neighbor comes online. As long as the OSPF network is stable, LSAs between OSPF neighbors don't occur. An LSA identifies the interfaces of all OSPF-enabled routers in an area, and provides information that enables OSPF-enabled routers to select the shortest path to a destination. All LSA exchanges between OSPF-enabled routers are authenticated.

When a network of OSPF routers comes online, the following steps occur:

1. When OSPF routers come online, they send out hello packets to find other OSPF routers on their network segment.
2. When they discover other routers on their network segment, they generally become adjacent. Adjacent routers can exchange routing updates. For more information, see ["Adjacency" on page 2148](#).
3. A DR and BDR are elected from the available routers using priority settings and router ID. See ["Designated router and backup router" on page 2149](#), and ["DR and BDR election issues" on page 2160](#).
4. Link state updates are sent between adjacent routers to map the topology of the OSPF area.
5. Once complete, the DR floods the network with the updates to ensure all OSPF routers in the area have the same OSPF route database. After the initial update, there are very few required updates if the network is stable.

How OSPF works on FortiGate devices

When a FortiGate interface is connected to an OSPF area, that unit can participate in OSPF communications. FortiGate devices use the OSPF hello protocol to acquire neighbors in an area. A neighbor is any router that's directly connected to the same area as the FortiGate and is ideally adjacent with a state of Full. After initial contact, the FortiGate exchanges hello packets with its OSPF neighbors regularly to confirm that the neighbors can be reached.

The number of routes that a FortiGate can learn through OSPF depends on the network topology. A single unit can support tens of thousands of routes if the OSPF network is configured properly.

External routes

OSPF is an internal routing protocol. OSPF external routes are routes where the destination is using a routing protocol other than OSPF. OSPF handles external routes by adjusting the cost of the route to include the cost of the other routing protocol. There are two methods of calculating this cost, which are used for OSPF external1 (E1) and OSPF external2 (E2).

OSPF E1

In OSPF E1, the destination is outside the OSPF domain. This requires a different metric to be used beyond the normal OSPF metrics. The new metric of a redistributed route is calculated by adding the external cost and the OSPF cost together.

OSPF E2

OSPF E2 is the default external type when routes are redistributed outside of OSPF. With OSPF E2, the metric of the redistributed route is equivalent to the external cost only, expressed as an OSPF cost. Dropping the OSPF portion can be useful in a number of situations, for example, on border routers that have no OSPF portion or where the OSPF routing cost is negligible compared to the external routing cost.

Comparing E1 and E2

The best way to understand OSPF E1 and E2 routes is to check routing tables on OSPF routers. If you look at the routes on an OSPF border router, the redistributed routes will have an associated cost that represents only the external route, as there is no OSPF cost to the route due to it already being on the edge of the OSPF domain. However, if you look at that same route on a different OSPF router inside the OSPF routing domain, it has a higher associated cost, essentially the external cost plus the cost over the OSPF domain to that border router. The border router uses OSPF E2, where the internal OSPF router uses OSPF E1 for the same route.

Viewing external routes

When you're trying to determine the costs for routes in your network to predict how traffic will be routed, you need to see the external OSPF routes and their associated costs. On a FortiGate, you can use the CLI to find this information.

To view external routes - CLI:

You can view the whole routing table using the `get router info routing-table all` command to see all of the routes, including the OSPF external routes. To view a shorter list, you can use the `get router info routing-table ospf` command. The letter at the left will be either E1 or E2 for external OSPF routes. The output looks similar to the following, depending on what routes are in the routing table:

```
FGT620B# get router info routing-table all
Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
E1 - OSPF external type 1, E2 - OSPF external type 2
i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
* - candidate default

O*E2   0.0.0.0/0 [110/10] via 10.1.1.3, tunnel_wan2, 00:02:11
O      10.0.0.1/32 [110/300] via 10.1.1.3, tunnel_wan2, 00:02:11
S      0.0.0.0/0 [10/0] via 192.168.183.254, port2
S      1.0.0.0/8 [10/0] via 192.168.183.254, port2
```

Link state database and route updates

OSPF is based on links. The links between adjacent neighbor routers allow updates to be passed along the network. Network links allow the DR to flood the area with link state database (LSDB) updates. External links allow the OSPF area to connect to destinations outside the OSPF autonomous system. Information about these links is passed throughout the OSPF network as link state updates.

The LSDB contains the information that defines the complete OSPF area, but the LSDB isn't the routing table. It contains the information from all the link state updates passed along the network. When there are no more changes required and the network is stable, the LSDB on each router in the network is the same. The DR floods the LSDB to the area to ensure that each router has the same LSDB.

To calculate the best route (shortest path) to a destination, the FortiGate applies the Shortest Path First (SPF) algorithm, based on Dijkstra's algorithm, to the accumulated link state information. OSPF uses relative path cost metric for choosing the best route. The path cost can be any metric, but it's typically the bandwidth of the path, which is how fast traffic will get from one point to another.

The path cost, similar to distance for RIP, imposes a penalty on the outgoing direction of a FortiGate interface. The path cost of a route is calculated by adding all of the costs associated with the outgoing interfaces along the path to the destination. The lowest overall path cost indicates the best route, and generally the fastest route. Some brands of OSPF routers, such as Cisco, implement cost as a direct result of bandwidth between the routers. Generally this is a good cost metric because larger bandwidth means more traffic can travel without slowing down. To achieve this type of cost metric on FortiGate devices, you need to set the cost for each interface manually in the CLI.



The inter-area routes may not be calculated when a Cisco type ABR has no fully adjacent neighbor in the backbone area. In this situation, the router considers summary-LSAs from all Actively summary-LSAs from all Actively Attached areas ([RFC 3509](#)).

The FortiGate dynamically updates its routing table based on the results of the SPF calculation to ensure that an OSPF packet will be routed using the shortest path to its destination. Depending on the network topology, the entries in the FortiGate routing table may include:

- The addresses of networks in the local OSPF area (to which packets are sent directly)
- Routes to OSPF area border routers (to which packets destined for another area are sent)
- Routes to area boundary routers, if the network contains OSPF areas and non-OSPF domains, which reside on the OSPF network backbone and are configured to forward packets to destinations outside the OSPF AS.

OSPF route updates

Once the OSPF domain is established, there should be few updates required on a stable network. When updates occur and a decision is required concerning a new route, this is the general procedure.

Our router gets a new route and needs to decide if it should go in the routing table.

The router has an up-to-date LSDB of the entire area, containing information about each router, the next hop to it, and most importantly the cost to get there.

Our router turns the LSDB into an SPF tree using Dijkstra's algorithm. It doesn't matter if there's more than one path to a router on the network, the SPF tree only cares about the shortest path to that router.

Once the SPF tree has been created and shows the shortest paths to all the OSPF routers on the network, the work is done. If the new route is the best route, it'll be part of that tree. If it's not the shortest route, it won't be included in the LSDB.

If there has been a change from the initial LSDB to the new SPF tree, a link state update will be sent out to let the other routers know about the change so they can also update their LSDBs. This is vital since all routers on the OSPF area must have the same LSDB.

If there was no change between the LSDB and the SPF tree, no action is taken.

OSPF packets

Every OSPF packet starts with a standard 24-byte header, and another 24 bytes of information or more. The header contains all the information necessary to determine whether the packet should be accepted for further processing.

OSPF packet

1-byte Version field	1-byte Type field	2-byte Packet length	3-byte Router ID
4-byte Area ID	2-byte Checksum	2-byte Auth Type	8-byte Authentication
4-byte Network Mask	2-byte Hello interval	1-byte Options field	1-byte Router Priority
4-byte Dead Router interval	4-byte DR field	4-byte BDR field	4-byte Neighbor ID

The following descriptions summarize the OSPF packet header fields:

Version field: The OSPF version number. This specification documents version 2 of the protocol.

Type field: There are 5 OSPF packet types. From one to five, respectively, they are Hello, Database Description, Link State Request, Link State Update, and Link State Acknowledgment.

Packet length: The length of the OSPF protocol packet, in bytes. This length includes the standard OSPF 24-byte header, so all OSPF packets are at 24-bytes long.

Router ID: The Router ID of the packet's source.

Area ID: A 32-bit number identifying the area that this packet belongs to. All OSPF packets are associated with a single area. Most travel a single hop only. Packets travelling over a virtual link are labelled with the backbone Area ID of 0.0.0.0.

Checksum: The standard IP checksum of the entire contents of the packet, starting with the OSPF packet header but excluding the 64-bit authentication field. This checksum is calculated as the 16-bit one's complement of the one's complement sum of all the 16-bit words in the packet, excepting the authentication field. If the packet's length isn't an integral number of 16-bit words, the packet is padded with a byte of zero before checksumming. The checksum is considered to be part of the packet authentication procedure. For some authentication types, the checksum calculation is omitted.

Auth Type: Identifies the authentication procedure to be used for the packet. Authentication types include Null authentication (0), Simple password (1), Cryptographic authentication (2), and all others are reserved for future use.

Authentication: A 64-bit field for use by the authentication scheme. When AuType indicates no authentication is being used, the authentication field isn't checked and can be any value. When AuType is set to 2 (cryptographic authentication), the 64-bit authentication field is split into the following four fields: Zero field, Key ID field, Authentication data length field, and Cryptographic sequence field.

The Key ID field indicates the key and algorithm used to create the message digest appended to the packet. The Authentication data length field indicates how many bytes long the message digest is. The Cryptographic sequence field is a non-decreasing number that is set when the packet is received and authenticated to prevent replay attacks.

Network Mask: The subnet where this packet is valid.

Hello interval: The period of time between sending out hello packets. For more information, see ["Hello and dead intervals" on page 2152](#).

Options field: The OSPF protocol defines several optional capabilities. A router indicates the optional capabilities that it supports in its OSPF hello packets, database description packets and in its LSAs. This enables routers supporting a mix of optional capabilities to coexist in a single AS.

Router priority: The priority, between 0 and 255, that determines which routers become the DR and BDR. For more information, see ["Designated router and backup router" on page 2149](#).

Dead router interval: The period of time when there's no response from a router before it's declared dead. For more information, see ["Hello and dead intervals" on page 2152](#).

DR and BDR fields: The DR and BDR fields each list the router that fills that role on this network, generally the routers with the highest priorities. For more information, see ["Designated router and backup router" on page 2149](#).

Neighbor ID: The ID number of a neighboring router. This ID is used to discover new routers and respond to them.

OSPF and VRFs

You can configure OSPF support for multiple virtual routing and forwarding (VRFs) on a FortiGate.

To add VRFs for interfaces - CLI:

```
config system interface
  edit <name>
    set vrf <VRF ID>
  next
end
```

where `vrf` is a value of 0 to 31.

Troubleshooting OSPF

As with other dynamic routing protocols, OSPF has some issues that may need troubleshooting from time to time. For basic troubleshooting, see the *Troubleshooting Handbook*.

Clearing OSPF routes from the routing table

If you think the wrong route has been added to your routing table and you want to check it out, you first have to remove that route from your table before seeing if it's added back in or not. You can clear all or some OSPF neighbor connections (sessions) using the `execute router clear ospf` CLI command. The `exec`

`router clear` command is much more limiting for OSPF than it is for BGP. For more information, see ["BGP" on page 2187](#).

For example, if you have routes in the OSPF routing table and you want to clear the specific route to IP address 10.10.10.1, you'll have to clear all the OSPF entries. Enter the following CLI command:

```
execute router clear ospf process
```

Checking the state of OSPF neighbors

In OSPF, each router sends out link state advertisements to find other routers on its network segment and to create adjacencies with some of those routers. This is important because routing updates are only passed between adjacent routers. If two routers you believe to be adjacent are not, that can be the source of routing failures.

To identify this problem, you need to check the state of the OSPF neighbors of the FortiGate. Use the `get router info ospf neighbor all` CLI command to see all the neighbors for the FortiGate. You'll see output in the form of the following:

```
FGT1 # get router info ospf neighbor
OSPF process 0:
Neighbor ID  Pri  State      Dead Time Address Interface
10.0.0.2      1   Full/ -   00:00:39 10.1.1.2 tunnel_wan1
10.0.0.2      1   Full/ -   00:00:34 10.1.1.4 tunnel_wan2
```

The important information here is the `State` column. Any neighbors that are not adjacent to the FortiGate are reported in this column as something other than `Full`. If the state is `Down`, that router is offline.

Passive interface problems

A passive OSPF interface doesn't send out any updates. This means it can't be a DR, BDR, or an area border router among other things. It depends on other neighbor routers to update its link state table.

Passive interfaces can cause problems when they're not receiving the routing updates you expect from their neighbors. This results in the passive OSPF interface on the FortiGate having an incomplete or out-of-date link state database, and it won't be able to properly route its traffic. It's possible that the passive interface is causing a hole in the network where no routers are passing updates to each other, however, this is a rare situation.

If a passive interface is causing problems, there are simple methods to determine it's the cause. The easiest method is to make it an active interface, and if the issues disappear, then that was the cause. Another method is to examine the OSPF routing table and related information to see if it's incomplete compared to other neighbor routers. If this is the case, you can clear the routing table, reset the device, and allow it to repopulate the table.

If you can't make the interface active for some reason, you have to change your network to fix the hole by adding more routers, or changing the relationship between the passive router's neighbors to provide better coverage.

Timer problems

A timer mismatch is when two routers have different values set for the same timer. For example, if one router declares a router dead after 45 seconds and another waits for 4 minutes, that difference in time results in the two routers being out of synch for that period of time. One will still see the offline router as being online.

The easiest method to check the timers is to check the configuration on each router. Another method is to sniff some packets, and read the timer values in the packets themselves from different routers. Each packet contains the hello interval and dead interval periods, so you can compare them easily enough.

BFD

Bidirectional Forwarding Detection (BFD) is a protocol that you can use to quickly locate hardware failures in the network. Routers running BFD communicate with each other and if a timer runs out on a connection then that router is declared down. BFD then communicates this information to the routing protocol and the routing information is updated. For more information about BFD, see ["BFD" on page 2106](#).

Authentication issues

OSPF has a number of authentication methods you can choose from. You may encounter problems with routers not authenticating as you expect. This will likely appear simply as one or more routers that have a blind spot in their routing and they won't acknowledge a router. This can be a problem if that router connects areas to the backbone, as it'll appear to be offline and unusable.

To confirm this is the issue, the easiest method is to turn off authentication on the neighboring routers. With no authentication between any routers, everything should flow normally.

Another method to confirm that authentication is the problem is to sniff packets and look at their contents. The authentication type and password are right in the packets which makes it easy to confirm they are what you expect during real time. It's possible that one or more routers isn't configured as you expect and may be using the wrong authentication. This method is especially useful if there are a group of routers with these problems since it may be only one router causing the problem that's seen in multiple routers.

Once you have confirmed the problem is related to authentication, you can decide how to handle it. You can turn off authentication and take your time to determine how to get your preferred authentication type back online. You can try another type of authentication, such as text instead of md5, which may have more success and still provide some level of protection. The important part is that once you confirm the problem, you can decide how to fix it properly.

DR and BDR election issues

You can force a particular router to become the DR and BDR by setting its priorities higher than any other OSPF routers in the area. This is a good idea when those routers have more resources to handle the traffic and extra work of the DR and BDR roles, since not all routers may be able to handle all of that traffic.

However, if you set all the other routers so they don't have a chance at being elected (give them a priority of 0), you can run into problems if the DR and BDR go offline. The good part is that you'll have some warning generally as the DR goes offline and the BDR is promoted to the DR position. However, if the network segment with both the DR and BDR goes down, your network won't have a way to send hello packets, send updates, or perform the other tasks that the DR performs.

The solution to this is to always allow routers to have a chance to be promoted, even if you set their priority to 1. In that case, they'll be the last choice but if there are no other candidates, you want that router to become the DR. Most networks will have already alerted you to the equipment problems, so this will be a temporary measure to keep the network traffic moving until you can find and fix the problem and get the real DR back online.

Basic OSPF example

This example sets up an OSPF network at a small office. There are 3 routers, all running OSPFv2. The border router connects to a BGP network.

All three routers in this example are FortiGate devices. Router1 will be the designated router (DR) and Router2 will be the backup designated router (BDR) due to their priorities. Router3 won't be considered for either the DR or

BDR elections. Instead, Router3 is the Autonomous System Border Router (ASBR) routing all traffic to the ISP's BGP router on its way to the Internet.

Router2 has a modem connected that provides dialup access to the Internet as well, at a reduced bandwidth. This is a PPPoE connection to a DSL modem. This provides an alternate route to the Internet if the other route goes down. The DSL connection is slow and is charged by the amount of traffic. For these reasons, OSPF will highly favor Router3's Internet access.

The DSL connection connects to an OSPF network with the ISP, so no redistribution of routes is required. However, the ISP network does have to be added to that router's configuration.

Network layout and assumptions

There are three FortiGate devices acting as OSPFv2 routers on the network: Router1, Router2, and Router3. Router1 will be the DR, and Router 2 the BDR. Router3 is the ASBR that connects to the external ISP router running BGP. Router2 has a PPPoE DSL connection that can access the Internet.

The head office network is connected to Router1 and Router2 on the 10.11.101.0 subnet.

Router1 and Router3 are connected over the 10.11.103.0 subnet.


Router2 and Router3 are connected over the 10.11.102.0 subnet.

The following table lists the router, interface, address, and role it's assigned.

Routers, interfaces, and IP addresses for the basic OSPF example network

Router name	Interface	IP address	Interface is connected to:
Router1 (DR)	Internal (port1)	10.11.101.1	Head office network and Router2
	External (port2)	10.11.102.1	Router3
Router2 (BDR)	Internal (port1)	10.11.101.2	Head office network and Router1
	External (port2)	10.11.103.2	Router3
	DSL (port3)	10.12.101.2	PPPoE DSL access
	Internal1 (port1)	10.11.102.3	Router1
Router3 (ASBR)	Internal2 (port2)	10.11.103.3	Router2
	External (port3)	172.20.120.3	ISP's BGP network

Basic OSPF network topology



Note that other subnets can be added to the internal interfaces without changing the configuration.

Assumptions

- The FortiGate devices used in this example have interfaces named port1, port2, and port3.
- All FortiGate devices in this example have factory default configuration with FortiOS 4.0 MR2 firmware installed and are in NAT/Route operation mode.
- Basic firewalls are in place to allow unfiltered traffic between all connected interfaces in both directions.
- This OSPF network is not connected to any other OSPF networks.
- Both Internet connections are always available.
- The modem connection is very slow and expensive.
- Other devices may be on the network, but do not affect this basic configuration.
- Router3 is responsible for redistributing all routes into and out of the OSPF AS.

Configuring the FortiGate devices

Each FortiGate needs the interfaces and basic system information, such as hostname, configured.

Configuring Router1

Router1 has two interfaces connected to the network: internal (port1) and external (port2). Its host name must be changed to Router1.

To configure Router1 interfaces - GUI:

1. Go to **System > Dashboard > Status**.
2. Beside the host name, select **Change**.
3. Enter a hostname of `Router1` and select **OK**.
4. Go to **Network > Interfaces**, edit port1, set the following information, and select **OK**.

Alias	internal
IP/Network Mask	10.11.101.1/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Head office and Router2
Administrative Status	Up

5. Edit port2, set the following information and select **OK**.

Alias	External
IP/Network Mask	10.11.102.1/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router3
Administrative Status	Up

Configuring Router2

Router2 configuration is the same as Router1, except Router2 also has the DSL interface to configure.

The DSL interface is configured with a username of “user1” and a password of “ospf_example”. The default gateway is retrieved from the ISP and the defaults are used for the rest of the PPPoE settings.

To configure Router2 interfaces - GUI:

1. Go to **System > Dashboard > Status**.
2. Beside the host name, select **Change**.
3. Enter a hostname of `Router2` and select **OK**.
4. Go to **Network > Interfaces**, edit port1, set the following information, and select **OK**.

Alias	internal
IP/Network Mask	10.11.101.2/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Head office and Router1

Administrative Status	Up
------------------------------	----

5. Edit port2, set the following information and select **OK**.

Alias	External
IP/Network Mask	10.11.103.2/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router3
Administrative Status	Up

6. Edit DSL (port3), set the following information and select **OK**.

Alias	DSL
Addressing Mode	PPPoE
Username	user1
Password	ospf_example
Unnumbered IP	10.12.101.2/255.255.255.0
Retrieve default gateway from server	Enable
Administrative Access	HTTPS SSH PING
Description	DSL
Administrative Status	Up

Configuring Router3

Router3 is similar to Router1 and Router2 configurations. The main difference is the External (port3) interface connected to the ISP BGP network, which has no administration access enabled, for security reasons.

To configure Router3 interfaces - GUI:

1. Go to **System > Status > Dashboard**.
2. Next to hostname, select **Change**.
3. Enter a hostname of `Router3` and select **OK**.
4. Go to **Network > Interfaces**, edit port1, set the following information, and select **OK**.

Alias	internal
--------------	----------

IP/Network Mask	10.11.102.3/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router1
Administrative Status	Up

5. Edit port2, set the following information and select **OK**.

Alias	Internal2
IP/Network Mask	10.11.103.3/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router2
Administrative Status	Up

6. Edit port3, set the following information and select **OK**.

Alias	External
IP/Network Mask	172.20.120.3/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	ISP BGP
Administrative Status	Up

Configuring OSPF on the FortiGate devices

With the interfaces configured, now the FortiGate devices can be configured for OSPF on those interfaces. All routers are part of the backbone 0.0.0.0 area, so no inter-area communications are needed.

For a simple configuration, there will be no authentication, no graceful restart or other advanced features, and timers will be left at their defaults. Also, the costs for all interfaces will be left at 10, except for the modem and ISP interfaces where cost will be used to load balance traffic. Nearly all advanced features of OSPF are only available from the CLI.

The network that's defined covers all the subnets used in this example - 10.11.101.0, 10.11.102.0, and 10.11.103.0. All routes for these subnets will be advertised. If there are other interfaces on the FortiGate devices that you don't want included in the OSPF routes, ensure those interfaces use a different subnet outside of the 10.11.0.0 network. If you want all interfaces to be advertised you can use an OSPF network of 0.0.0.0 .

Each router will configure:

- Router ID
- Area

- Network
- Two or three interfaces depending on the router
- Priority for DR (Router1) and BDR (Router2)
- Redistribute for ASBR (Router3)

Configuring OSPF on Router1

Router1 has a very high priority to ensure it becomes the DR for this area. Also Router1 has the lowest IP address to help ensure it will win in case there's a tie at some point. Otherwise, it's a standard OSPF configuration. Setting the priority can only be done in the CLI, and it's for a specific OSPF interface.

To configure OSPF on Router1 - GUI:

1. Go to **Router > Dynamic > OSPF**.
2. Set **Router ID** to `10.11.101.1` and select **Apply**.
3. In **Areas**, select **Create New**, set the following information, and select **OK**.

Area	0.0.0.0
Type	Regular
Authentication	none

4. In **Networks**, select **Create New**, set the following information, and select **OK**.

IP/Netmask	10.11.0.0/255.255.0.0
Area	0.0.0.0

5. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Router1-Internal-DR
Interface	port1 (Internal)
IP	0.0.0.0
Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	40

6. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Router1-External
-------------	------------------

Interface	port2 (External)
IP	0.0.0.0
Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	40

7. Using the CLI, enter the following commands to set the priority for the Router1-Internal OSPF interface to maximum, ensuring this interface becomes the DR:

```
config router ospf
  config ospf-interface
    edit Router1-Internal-DR
      set priority 255
    next
  end
```

To configure OSPF on Router1 - CLI:

```
config router ospf
  set router-id 10.11.101.1
  config area
    edit 0.0.0.0
    next
  end
  config network
    edit 1
      set prefix 10.11.0.0/255.255.255.0
    next
  end
  config ospf-interface
    edit "Router1-Internal"
      set interface "port1"
      set priority 255
    next
    edit "Router1-External"
      set interface "port2"
    next
  end
end
```

Configuring OSPF on Router2

Router2 has a high priority to ensure it becomes the BDR for this area and configures the DSL interface slightly differently. Assume this will be a slower connection resulting in the need for longer timers and a higher cost for this route.

Otherwise, it is a standard OSPF configuration.

To configure OSPF on Router2 - GUI:

1. Go to **Router > Dynamic > OSPF**.
2. Set **Router ID** to `10.11.101.2` and select **Apply**.
3. In **Areas**, select **Create New**, set the following information, and select **OK**.

Area	0.0.0.0
Type	Regular
Authentication	none

4. In **Networks**, select **Create New**, set the following information, and select **OK**.

IP/Netmask	10.11.0.0/255.255.0.0
Area	0.0.0.0

5. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Router2-Internal
Interface	port1 (Internal)
IP	0.0.0.0
Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	

6. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Router2-External
Interface	port2 (External)
IP	0.0.0.0
Authentication	none
Timers (seconds)	
Hello Interval	10
Dead Interval	40

7. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Router2-DSL
Interface	port3 (DSL)
IP	0.0.0.0
Authentication	none
Cost	50
Timers (seconds)	
Hello Interval	20
Dead Interval	80

8. Using the CLI, enter the following commands to set the priority for the Router2-Internal OSPF interface to ensure this interface will become the BDR:

```
config router ospf
  config ospf-interface
    edit Router2-Internal
      set priority 250
    next
  end
```

To configure OSPF on Router2 - CLI:

```
config router ospf
  set router-id 10.11.101.2
  config area
    edit 0.0.0.0
    next
  end
  config network
    edit 1
      set prefix 10.11.0.0/255.255.0.0
    next
  end
  config ospf-interface
    edit "Router2-Internal"
      set interface "port1"
      set priority 255
    next
    edit "Router2-External"
      set interface "port2"
    next
    edit "Router2-DSL"
      set interface "port3"
      set cost 50
    next
  end
end
```


Configuring OSPF on Router3

Router3 is more complex than the other two routers. The interfaces are straightforward, but this router has to import and export routes between OSPF and BGP. That requirement makes Router3 an ASBR. Also, Router3 needs a lower cost on its route to encourage all traffic to the Internet to route through it.

In the advanced OSPF options, redistribute is enabled for Router3. It allows different types of routes, learned outside of OSPF, to be used in OSPF. Different metrics are assigned to these other types of routes to make them more or less preferred to regular OSPF routes.

To configure OSPF on Router3 - GUI:

1. Go to **Router > Dynamic > OSPF**.
2. Set **Router ID** to `10.11.101.2` and select **Apply**.
3. Expand **Advanced Options**.
4. In **Redistribute**, set the following information, and select **OK**.

Route type	Redistribute	Metric
Connected	Enable	15
Static	Enable	15
RIP	Disable	n/a
BGP	Enable	5

5. In **Areas**, select **Create New**, set the following information, and select **OK**.

Area	0.0.0.0
Type	Regular
Authentication	none

6. In **Networks**, select **Create New**, set the following information, and select **OK**.

IP/Netmask	10.11.0.0/255.255.0.0
Area	0.0.0.0

7. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Router3-Internal
Interface	port1 (Internal)
IP	0.0.0.0
Authentication	none

Timers (seconds)**Hello Interval** 10**Dead Interval** 40

8. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name Router3-Internal2**Interface** port2 (Internal2)**IP** 0.0.0.0**Authentication** none**Timers (seconds)****Hello Interval** 10**Dead Interval** 40

9. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name Router3-ISP-BGP**Interface** port3 (ISP-BGP)**IP** 0.0.0.0**Authentication** none**Cost** 2**Timers (seconds)****Hello Interval** 20**Dead Interval** 80

10. Using the CLI, enter the following commands to set the priority for the Router3-Internal OSPF interface to ensure this interface will become the BDR:

```
config router ospf
  config ospf-interface
    edit Router3-Internal
      set priority 250
    next
  end
```

To configure OSPF on Router3 - CLI:

```
config router ospf
  set router-id 10.11.102.3
  config area
    edit 0.0.0.0
    next
  end
  config network
    edit 1
      set prefix 10.11.0.0/255.255.255.0
    next
    edit 2
      set prefix 172.20.120.0/255.255.255.0
    next
  end
  config ospf-interface
    edit "Router3-Internal"
      set interface "port1"
      set priority 255
    next
    edit "Router3-External"
      set interface "port2"
    next
    edit "Router3-ISP-BGP"
      set interface "port3"
      set cost 2
    next
  end
end
```

Configuring other networking devices

The other networking devices required in this configuration are on the two ISP networks, the BGP network for the main Internet connection, and the DSL backup connection.

In both cases, the ISPs need to be notified about the OSPF network settings including router IP addresses, timer settings, and so on. The ISP will use this information to configure its routers that connect to this OSPF network.

Testing network configuration

Testing the network configuration involves two parts: testing the network connectivity and testing the OSPF routing.

To test the network connectivity, use ping, traceroute, and other network tools.

To test the OSPF routing in this example, refer to the troubleshooting outlined in ["Troubleshooting OSPF" on page 2158](#).

Advanced inter-area OSPF example

This example sets up an OSPF network at a large office. There are three areas, each with two routers. Typically OSPF areas wouldn't be this small, and if they were, the areas would be combined into one larger area. However, the stub area services the accounting department whose members are very sensitive about their network and don't want their network information broadcasted through the rest of the company. The backbone area contains

the bulk of the company's network devices. The regular area was established for various reasons, such as hosting the company servers in a separate area with extra security.

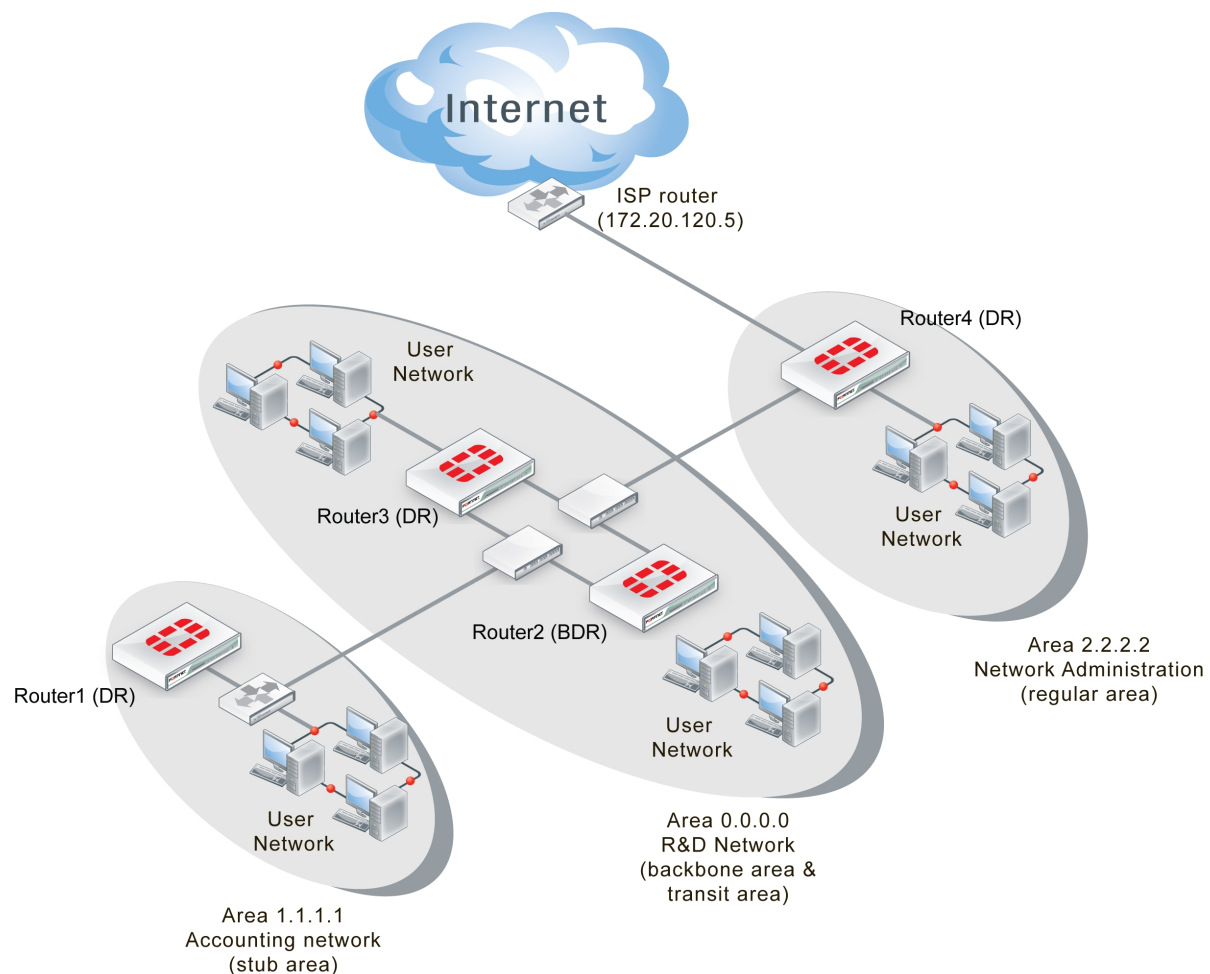
One area is a small stub area that has no independent Internet connection, and has only one connection to the backbone area. That connection between the stub area and the backbone area is only through a default route. No routes outside the stub area are advertised into that area. Another area is the backbone, which is connected to the other two areas. The third area has the Internet connection, and all traffic to and from the Internet must use that area's connection. If that traffic comes from the stub area, then that traffic is treating the backbone like a transit area that only uses it to get to another area.

In the stub area, a subnet of computers is running the RIP routing protocol and those routes must be redistributed into the OSPF areas.

Network layout and assumptions

There are four FortiGate devices in this network topology, which are acting as OSPF routers:

Advanced inter-area OSPF network topology



Area 1.1.1.1 is a stub area with one FortiGate OSPF router called Router1 (DR). Its only access outside of that area is a default route to the backbone area, which is how it accesses the Internet. Traffic must go from the stub area, through the backbone, to the third area to reach the Internet. The backbone area in this configuration is

called a transit area. Also, in area 1.1.1.1 there is a RIP router that will be providing routes to the OSPF area through redistribution.

Area 0.0.0.0 is the backbone area and has two FortiGate device routers named Router2 (BDR) and Router3 (DR).

Area 2.2.2.2 is a regular area that has an Internet connection accessed by both the other two OSPF areas. There is only one FortiGate device router in this area called Router4 (DR). This area is more secure and requires MD5 authentication by routers.

All areas have user networks connected but they're not important for configuring the network layout for this example.

Internal interfaces are connected to internal user networks only. External1 interfaces are connected to the 10.11.110.0 network, joining Area 1.1.1.1 and Area 0.0.0.0.

External2 interfaces are connected to the 10.11.111.0 network, joining Area 0.0.0.0 and Area 2.2.2.2. The ISP interface is called ISP.

Routers, areas, interfaces, and IP addresses for advanced OSPF network

Router name	Area number and type	Interface	IP address
Router1 (DR)	1.1.1.1 - stub area (Accounting)	port1 (internal)	10.11.101.1
		port2 (external1)	10.11.110.1
Router2 (BDR)	0.0.0.0 - backbone area (R&D Network)	port1 (internal)	10.11.102.2
		port2 (external1)	10.11.110.2
		port3 (external2)	10.11.111.2
Router3 (DR)	0.0.0.0 - backbone area (R&D Network)	port1 (internal)	10.11.103.3
		port2 (external1)	10.11.110.3
		port3 (external2)	10.11.111.3
Router4 (DR)	2.2.2.2 - regular area (Network Admin)	port1 (internal)	10.11.104.4
		port2 (external2)	10.11.111.4
		port3 (ISP)	172.20.120.4

Note that other subnets can be added to the internal interfaces without changing the configuration.

Assumptions

- The FortiGate devices used in this example have interfaces named port1, port2, and port3.
- All FortiGate devices in this example have factory default configuration with FortiOS 4.0 MR2 firmware installed and are in NAT/Route operation mode.
- During configuration, if settings are not directly referred to, they will be left at the default settings.

- Basic firewalls are in place to allow unfiltered traffic between all connected interfaces in both directions.
- This OSPF network is not connected to any other OSPF areas outside of this example.
- The Internet connection is always available.
- Other devices may be on the network but do not affect this configuration.

Configuring the FortiGate devices

This section configures the basic settings on the FortiGate devices to be OSPF routers. These configurations include multiple interface settings and the hostname.

There are four FortiGate devices in this example. The two devices in the backbone area can be configured exactly the same except for IP addresses, so only the Router3 (the DR) configuration will be given, with notes indicating Router2's (the BDR) IP addresses.

Configuring the FortiGate devices includes:

- [Configuring Router1](#)
- [Configuring Router2](#)
- [Configuring Router3](#)
- [Configuring Router4](#)

Configuring Router1

Router1 is part of the Accounting network stub area (1.1.1.1).

To configure Router1 interfaces - GUI:

1. Go to **System > Dashboard > Status**.
2. Next to hostname, select **Change**.
3. Enter a hostname of `Router1` and select **OK**.
4. Go to **Network > Interfaces** edit port1, set the following information, and select **OK**.

Alias	internal
IP/Network Mask	10.11.101.1/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Accounting network
Administrative Status	Up

5. Edit port2, set the following information and select **OK**.

Alias	External1
IP/Network Mask	10.11.110.1/255.255.255.0
Administrative Access	HTTPS SSH PING

Description	Backbone network and Internet
Administrative Status	Up

Configuring Router2

Router2 is part of the R&D network backbone area (0.0.0.0). Router2 and Router3 are in this area. They provide a redundant connection between area 1.1.1.1 and area 2.2.2.2.

Router2 has three interfaces configured: one to the internal network and two to Router3 for redundancy.

To configure Router2 interfaces - GUI:

1. Go to **System > Dashboard > Status**.
2. Next to hostname, select **Change**.
3. Enter a hostname of `Router2` and select **OK**.
4. Go to **Network > Interfaces**, edit port1 (internal), set the following information, and select **OK**.

Alias	internal
IP/Network Mask	10.11.102.2/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Internal RnD network
Administrative Status	Up

5. Edit port2 (external1), set the following information and select **OK**.

Alias	external1
IP/Network Mask	10.11.110.2/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router3 first connection
Administrative Status	Up

6. Edit port3 (external2), set the following information and select **OK**.

Alias	external2
IP/Network Mask	10.11.111.2/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router3 second connection

Administrative Status	Up
------------------------------	----

Configuring Router3

Router3 is part of the R&D network backbone area (0.0.0.0). Router2 and Router3 are in this area. They provide a redundant connection between area 1.1.1.1 and area 2.2.2.2.

To configure Router3 interfaces - GUI:

1. Go to **System > Dashboard > Status**.
2. Next to hostname, select **Change**.
3. Enter a hostname of `Router3` and select **OK**.
4. Go to **Network > Interfaces**, edit port1 (internal), set the following information, and select **OK**.

Alias	internal
IP/Network Mask	10.11.103.3/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Internal RnD network
Administrative Status	Up

5. Edit port2 (external1), set the following information and select **OK**.

Alias	external1
IP/Network Mask	10.11.110.3/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router2 first connection
Administrative Status	Up

6. Edit port3 (external2), set the following information and select **OK**.

Alias	external2
IP/Network Mask	10.11.111.3/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Router2 second connection
Administrative Status	Up

Configuring Router4

Router4 is part of the Network Administration regular area (2.2.2.2). This area provides Internet access for both area 1.1.1.1 and the backbone area.

This section configures interfaces and hostname.

To configure Router4 interfaces - GUI:

1. Go to **System > Dashboard > Status**.
2. Next to hostname, select **Change**.
3. Enter a hostname of `Router4` and select **OK**.
4. Go to **Network > Interfaces**.
5. Edit port1 (internal).
6. Set the following information and select **OK**.

Alias	internal
IP/Network Mask	10.11.101.4/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Accounting network
Administrative Status	Up

7. Edit port2 (external2).
8. Set the following information and select **OK**.

Alias	external2
IP/Network Mask	10.11.110.4/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	Backbone and Accounting network
Administrative Status	Up

9. Edit port3 (ISP).
10. Set the following information and select **OK**.

Alias	ISP
IP/Network Mask	172.20.120.4/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	ISP and Internet

Administrative Status	Up
------------------------------	----

Configuring OSPF on the FortiGate devices

Three of the routers are designated routers (DR) and one is a backup DR (BDR). This is achieved through the lowest router ID numbers, or OSPF priority settings.

Also, each area needs to be configured as each respective type of area: stub, backbone, or regular. This affects how routes are advertised into the area.

To configure OSPF on Router1 - GUI:

1. Go to **Router > Dynamic > OSPF**.
2. Enter **10.11.101.1** for the **Router ID** and select **Apply**.
3. In **Areas**, select **Create New**, set the following information, and select **OK**.

Area	1.1.1.1
Type	Stub
Authentication	None

4. In **Networks**, select **Create New**, set the following information, and select **OK**.

IP/Netmask	10.11.101.0/255.255.255.0
Area	1.1.1.1

5. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Accounting
Interface	port1 (internal)
IP	10.11.101.1
Authentication	None

6. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Backbone1
Interface	port2 (external1)
IP	10.11.110.1
Authentication	None

To configure OSPF on Router2 - GUI:

1. Go to **Router > Dynamic > OSPF**.
2. Enter 10.11.102.2 for the **Router ID** and select **Apply**.
3. In **Areas**, select **Create New**, set the following information, and select **OK**.

Area	0.0.0.0
Type	Regular
Authentication	None

4. In **Networks**, select **Create New**, set the following information, and select **OK**.

IP/Netmask	10.11.102.2/255.255.255.0
Area	0.0.0.0

5. In **Networks**, select **Create New**, set the following information, and select **OK**.

IP/Netmask	10.11.110.2/255.255.255.0
Area	0.0.0.0

6. In **Networks**, select **Create New**, set the following information, and select **OK**.

IP/Netmask	10.11.111.2/255.255.255.0
Area	0.0.0.0

7. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	RnD network
Interface	port1 (internal)
IP	10.11.102.2
Authentication	None

8. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Backbone1
Interface	port2 (external1)
IP	10.11.110.2
Authentication	None

9. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Backbone2
Interface	port3 (external2)
IP	10.11.111.2
Authentication	None

To configure OSPF on Router3 - GUI:

1. Go to **Router > Dynamic > OSPF**.
2. Enter 10.11.103.3 for the **Router ID** and select **Apply**.
3. In **Areas**, select **Create New**, set the following information, and select **OK**.

Area	0.0.0.0
Type	Regular
Authentication	None

4. In **Networks**, select **Create New**, set the following information, and select **OK**.

IP/Netmask	10.11.102.3/255.255.255.0
Area	0.0.0.0

5. In **Networks**, select **Create New**, set the following information, and select **OK**.

IP/Netmask	10.11.110.3/255.255.255.0
Area	0.0.0.0

6. In **Networks**, select **Create New**, set the following information, and select **OK**.

IP/Netmask	10.11.111.3/255.255.255.0
Area	0.0.0.0

7. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	RnD network
Interface	port1 (internal)
IP	10.11.103.3
Authentication	None

8. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Backbone1
Interface	port2 (external1)
IP	10.11.110.3
Authentication	None

9. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Backbone2
Interface	port3 (external2)
IP	10.11.111.3
Authentication	None

To configure OSPF on Router4 - GUI:

1. Go to **Router > Dynamic > OSPF**.
2. Enter 10.11.104.4 for the **Router ID** and then select **Apply**.
3. In **Areas**, select **Create New**.
4. Set the following information and select **OK**.

Area	2.2.2.2
Type	Regular
Authentication	None

5. In **Networks**, select **Create New**, set the following information, and select **OK**.

IP/Netmask	10.11.104.0/255.255.255.0
Area	0.0.0.0

6. In **Networks**, select **Create New**, set the following information, and select **OK**.

IP/Netmask	10.11.111.0/255.255.255.0
Area	0.0.0.0

7. In **Networks**, select **Create New**, set the following information, and select **OK**.

IP/Netmask	172.20.120.0/255.255.255.0
Area	0.0.0.0

8. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Network Admin network
Interface	port1 (internal)
IP	10.11.104.4
Authentication	None

9. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	Backbone2
Interface	port2 (external2)
IP	10.11.111.4
Authentication	None

10. In **Interfaces**, select **Create New**, set the following information, and select **OK**.

Name	ISP
Interface	port3 (ISP)
IP	172.20.120.4
Authentication	None

Configuring other networking devices

All network devices on this network are running OSPF routing. The user networks (Accounting, R&D, and Network Administration) are part of one of the three areas.

The ISP needs to be notified of your network configuration for area 2.2.2.2. Your ISP won't advertise your areas externally as they're intended as internal areas. External areas have assigned unique numbers. The area numbers used in this example are similar to the 10.0.0.0 and 192.168.0.0 subnets used in internal networking.

Testing network configuration

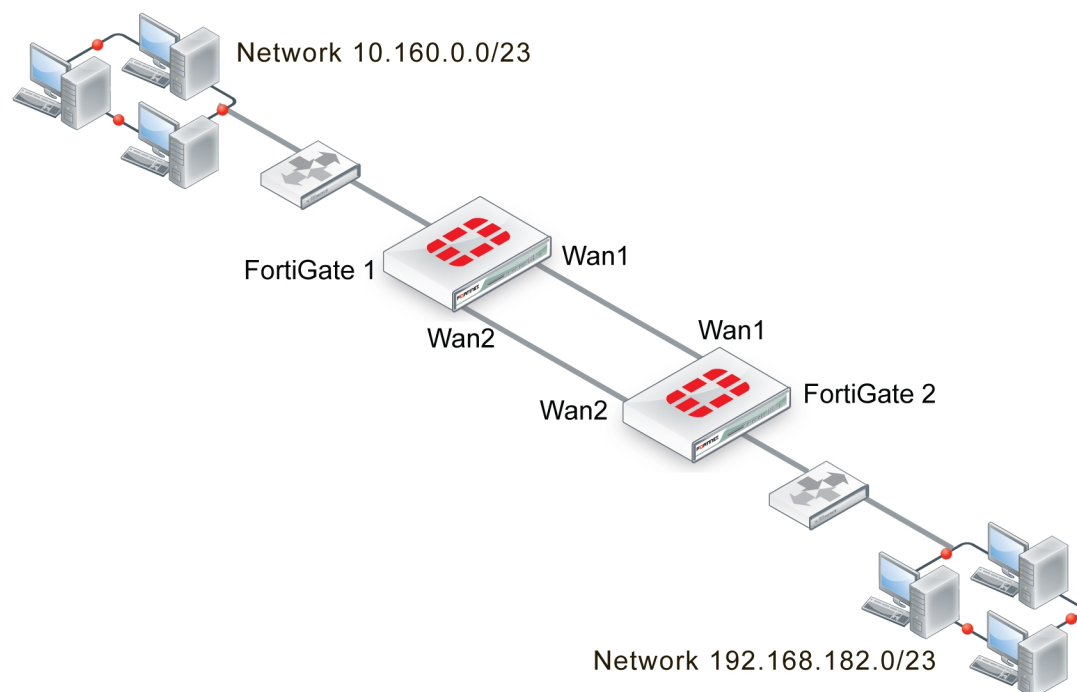
There are two main areas to test in this network configuration: network connectivity and OSPF routing.

To test network connectivity, see if computers on the Accounting or R&D networks can access the Internet. If you need troubleshooting network connectivity, see the *Troubleshooting Handbook*.

To test OSPF routing, check the routing tables on the FortiGate devices to ensure the expected OSPF routes are present. If you need help troubleshooting OSPF routing, see ["Troubleshooting OSPF" on page 2158](#).

Controlling redundant links by cost

In this scenario, two FortiGate devices have redundant links: one link between their WAN1 interfaces and one link between their WAN2 interfaces.



FortiGate 1 should learn the route to network 192.168.182.0 and FortiGate 2 should learn the route to network 10.160.0.0. Under normal conditions, they should learn these routes through the WAN1 link. The WAN2 link should be used only as a backup.

With the default settings, each FortiGate learns these routes from both WAN1 and WAN2.

FortiGate 1:

```
FGT1 # get router info ospf neighbor
OSPF process 0:
Neighbor ID Pri State Dead Time Address Interface
10.2.2.2 1 Full/Backup 00:00:33 10.182.0.187 wan1
10.2.2.2 1 Full/Backup 00:00:31 10.183.0.187 wan2
FGT1 # get router info routing-table ospf
O*E2 0.0.0.0/0 [110/10] via 10.183.0.187, wan2, 00:00:01
[110/10] via 10.182.0.187, wan1, 00:00:01
O 192.168.182.0/23 [110/20] via 10.183.0.187, wan2, 00:02:04
[110/20] via 10.182.0.187, wan1, 00:02:04
```

FortiGate 2:

```
FGT2 # get router info ospf neighbor
OSPF process 0:
Neighbor ID Pri State Dead Time Address Interface
10.1.1.1 1 Full/DR 00:00:38 10.182.0.57 wan1
10.1.1.1 1 Full/DR 00:00:38 10.183.0.57 wan2
FGT2 # get router info routing-table ospf
O 10.160.0.0/23 [110/20] via 10.183.0.57, wan2, 00:00:39
```

```
[110/20] via 10.182.0.57, wan1, 00:00:39
```

Adjusting the route costs

On both FortiGate devices, the cost of the route through WAN2 is adjusted higher so that this route will only be used if the route through WAN1 is unavailable. The default cost is 10. The WAN2 route will be changed to a cost of 200.

On both FortiGate devices:

```
config router ospf
  config ospf-interface
    edit "WAN2_higher_cost"
      set cost 200
      set interface "wan2"
    end
```

Now, both FortiGate devices use only the WAN1 route:

FortiGate 1:

```
FGT1 # get router info routing-table ospf
O*E2 0.0.0.0/0 [110/10] via 10.182.0.187, wan1, 00:00:40
O 192.168.182.0/23 [110/20] via 10.182.0.187, wan1, 00:00:40
```

FortiGate 2:

```
FGT2 # get router info routing-table ospf
O 10.160.0.0/23 [110/20] via 10.182.0.57, wan1, 00:09:37
```

LSDB check on FortiGate 1:

```
FGT1 # get router info ospf database router lsa
Router Link States (Area 0.0.0.0)
LS age: 81
Options: 0x2 (*|---|---|E|)
Flags: 0x0
LS Type: router-LSA
Link State ID: 10.1.1.1
Advertising Router: 10.1.1.1
LS Seq Number: 8000000b
Checksum: 0xe637
Length: 60
Number of Links: 3

Link connected to: Stub Network
(Link ID) Network/subnet number: 10.160.0.0
(Link Data) Network Mask: 255.255.254.0
Number of TOS metrics: 0
TOS 0 Metric: 10
Link connected to: a Transit Network
(Link ID) Designated Router address: 10.183.0.187
(Link Data) Router Interface address: 10.183.0.57
Number of TOS metrics: 0
TOS 0 Metric: 200

Link connected to: a Transit Network
```



```
(Link ID) Designated Router address: 10.182.0.57
(Link Data) Router Interface address: 10.182.0.57
Number of TOS metrics: 0
TOS 0 Metric: 10
```

```
LS age: 83
Options: 0x2 (*|---|E|)
Flags: 0x2 : ASBR
LS Type: router-LSA
Link State ID: 10.2.2.2
Advertising Router: 10.2.2.2
LS Seq Number: 8000000e
Checksum: 0xfc9b
Length: 60
  Number of Links: 3
```

```
Link connected to: Stub Network
(Link ID) Network/subnet number: 192.168.182.0
(Link Data) Network Mask: 255.255.254.0
Number of TOS metrics: 0
TOS 0 Metric: 10
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 10.183.0.187
(Link Data) Router Interface address: 10.183.0.187
Number of TOS metrics: 0
TOS 0 Metric: 200
```

```
Link connected to: a Transit Network
(Link ID) Designated Router address: 10.182.0.57
(Link Data) Router Interface address: 10.182.0.187
Number of TOS metrics: 0
TOS 0 Metric: 10
```

Verifying route redundancy

Bring down WAN1 and then check the routes on the two FortiGate devices.

FortiGate 1:

```
FGT1 # get router info routing-table ospf
FGT1 # get router info routing-table ospf
O*E2 0.0.0.0/0 [110/10] via 10.183.0.187, wan2, 00:00:06
O 192.168.182.0/23 [110/210] via 10.183.0.187, wan2, 00:00:06
```

FortiGate 2:

```
FGT2 # get router info routing-table ospf
O 10.160.0.0/23 [110/210] via 10.183.0.57, wan2, 00:00:14
```

The WAN2 interface is now in use on both units.

BGP

BGP contains two distinct subsets: internal BGP (iBGP) and external BGP (eBGP). iBGP is intended for use within your own networks. eBGP is used to connect many different networks together and is the main routing protocol for the Internet backbone. FortiGate devices support iBGP, and eBGP only for communities.

BGP was first used in 1989. The current version, BGP-4, was released in 1995 and is defined in [RFC 1771](#). That RFC has since been replaced by [RFC 4271](#). The main benefits of BGP-4 are classless inter-domain routing and aggregate routes. BGP is the only routing protocol to use TCP for a transport protocol. Other routing protocols use UDP.

BGP makes routing decisions based on path, network policies, and rulesets instead of the hop-count metric as RIP does, or cost-factor metrics as OSPF does.

BGP-4+ supports IPv6. It was introduced in [RFC 2858](#) and [RFC 2545](#).

BGP is the routing protocol used on the Internet. It was designed to replace the old Exterior Gateway Protocol (EGP) which had been around since 1982, and was very limited. BGP enabled more networks to take part in the Internet backbone to effectively decentralize it and make the Internet more robust, and less dependent on a single ISP or backbone network.

Parts and terminology of BGP

In a BGP network, there are some terms that need to be explained before going ahead. Some parts of BGP aren't explained here because they're common to other dynamic routing protocols. When determining your network topology, note that the number of available or supported routes isn't set by the configuration but depends on the available memory on the FortiGate. For more information about the parts of BGP that aren't listed here, see ["Dynamic routing" on page 2093](#).

BGP and IPv6

FortiGate devices support IPv6 over BGP using the same `config router bgp` CLI command as IPv4 but different subcommands.

The main CLI keywords have IPv6 equivalents that are identified by the "6" on the end of the keyword, such as `config network6` or `set allowas-in6`. For more information about IPv6 BGP keywords, see the [FortiOS CLI Reference](#).

IPv6 BGP commands include:

```
config router bgp
  set activate6 {enable | disable}
  set allowas-in6 <max_num_AS_integer>
  set allowas-in-enable6 {enable | disable}
  set as-override6 {enable | disable}
  set attribute-unchanged6 [as-path] [med] [next-hop]
  set capability-default-originate6 {enable | disable}
  set capability-graceful-restart6 {enable | disable}
  set capability-origf6 {both | none | receive | send}
  set default-originate-route-map6 <routemap_str>
  set distribute-list-in6 <access-list-name_str>
  set distribute-list-out6 <access-list-name_str>
  set filter-list-in6 <aspath-list-name_str>
  set filter-list-out6 <aspath-list-name_str>
  set maximum-prefix6 <prefix_integer>
  set maximum-prefix-threshold6 <percentage_integer>
```

```

set maximum-prefix-warning-only6 {enable | disable}
set next-hop-self6 {enable | disable}
set prefix-list-in6 <prefix-list-name_str>
set prefix-list-out6 <prefix-list-name_str>
set remove-private-as6 {enable | disable}
set route-map-in6 <route-map-name_str>
set route-map-out6 <route-map-name_str>
set route-reflector-client6 {enable | disable}
set route-server-client6 {enable | disable}
set send-community6 {both | disable | extended | standard}
set soft-reconfiguration6 {enable | disable}
set unsuppress-map6 <route-map-name_str>
config network6
config redistribute6
end

```

Role of routers in BGP networks

Dynamic routing has a number of different roles that routers can fill, such as those covered in . BGP has a number of custom roles that routers can fill. These include speaker routers, peer routers or neighbors, and route reflectors.

Speaker routers

Any router that's configured for BGP is considered a BGP speaker. This means that a speaker router advertises BGP routes to its peers.

Any routers on the network that aren't speaker routers aren't treated as BGP routers.

Peer routers or neighbors

In a BGP network, all neighboring BGP routers or peer routers are routers that are connected to a FortiGate. A FortiGate learns about all other routers through these peers.

You need to manually configure BGP peers on a FortiGate as neighbors. Otherwise, these routers won't be seen as peers, but simply as other routers on the network that don't support BGP. Optionally, you can use MD5 authentication to password-protect BGP sessions with those neighbors (see [RFC 2385](#)).

You can configure up to 1000 BGP neighbors on a FortiGate. You can clear all or some BGP neighbor connections (sessions), using the `execute router clear bgp` CLI command.

For example, if you have 10 routes in the BGP routing table and you want to clear the specific route to IP address 10.10.10.1, enter the following CLI command:

```
execute router clear bgp ip 10.10.10.1
```

To remove all routes for AS number 650001, enter the following CLI command:


```
execute router clear bgp as 650001
```

To remove route flap dampening information for the 10.10.0.0/16 subnet, enter the following CLI command:

```
execute router clear bgp dampening 10.10.0.0/16
```

In the following diagram, Router A is directly connected to five other routers in a network that contains 12 routers. These routers (the ones in the blue circle) are Router A's peers or neighbors.

Router A and its five peer routers



As a minimum, when configuring BGP neighbors, you must enter their IP address and the AS number (remote-as). This is all of the information the GUI allows you to enter for a neighbor.

The BGP commands related to neighbors are quite extensive and include:

```
config router bgp
config neighbor
  edit <neighbor_address_ipv4>
    set activate {enable | disable}
    set advertisement-interval <seconds_integer>
    set allowas-in <max_num_AS_integer>
    set allowas-in-enable {enable | disable}
    set as-override {enable | disable}
    set attribute-unchanged [as-path] [med] [next-hop]
    set bfd {enable | disable}
    set capability-default-originate {enable | disable}
    set capability-dynamic {enable | disable}
    set capability-graceful-restart {enable | disable}
    set capability-orf {both | none | receive | send}
    set capability-route-refresh {enable | disable}
    set connect-timer <seconds_integer>
    set description <text_str>
    set distribute-list-in <access-list-name_str>
    set distribute-list-out <access-list-name_str>
    set dont-capability-negotiate {enable | disable}
    set ebgp-enforce-multihop {enable | disable}
    set ebgp-multihop {enable | disable}
    set ebgp-multihop-ttl <seconds_integer>
    set filter-list-in <aspath-list-name_str>
    set filter-list-out <aspath-list-name_str>
    set holdtime-timer <seconds_integer>
    set interface <interface-name_str>
    set keep-alive-timer <seconds_integer>
    set maximum-prefix <prefix_integer>
    set maximum-prefix-threshold <percentage_integer>
    set maximum-prefix-warning-only {enable | disable}
    set next-hop-self {enable | disable}
    set passive {enable | disable}
    set password <string>
```

```
    set prefix-list-in <prefix-list-name_str>
    set prefix-list-out <prefix-list-name_str>
    set remote-as <id_integer>
    set remove-private-as {enable | disable}
    set retain-stale-time <seconds_integer>
    set route-map-in <routemap-name_str>
    set route-map-out <routemap-name_str>
    set route-reflector-client {enable | disable}
    set route-server-client {enable | disable}
    set send-community {both | disable | extended | standard}
    set shutdown {enable | disable}
    set soft-reconfiguration {enable | disable}
    set strict-capability-match {enable | disable}
    set unsuppress-map <route-map-name_str>
    set update-source <interface-name_str>
    set weight <weight_integer>
end
end
end
```

Route reflectors

Route reflectors (RR) in BGP concentrate route updates so other routers only need to talk to the RRs to get all of the updates. This results in smaller routing tables, fewer connections between routers, faster responses to network topology changes, and less administration bandwidth. BGP RRs are defined in [RFC 1966](#).


In a BGP RR configuration, the AS is divided into different clusters that each include client and reflector routers. The client routers supply the reflector routers with the client's route updates. The reflectors pass this information along to other RRs and border routers. Only the reflectors need to be configured, not the clients, because the clients find the closest reflector and communicate with it automatically. The reflectors communicate with each other as peers. A FortiGate can be configured as either reflectors or clients.

Since RRs are processing more than the client routers, the reflectors should have more resources to handle the extra workload.

Smaller networks running BGP typically do not require RRs. However, RRs are a useful feature for large companies, where their AS may include 100 routers or more. For example, a full mesh 20 router configuration within an AS, there would have to be 190 unique BGP sessions just for routing updates within the AS. The number of sessions jumps to 435 sessions for just 30 routers, or 4950 sessions for 100 routers. Based on these numbers, updating this many sessions will quickly consume the limited bandwidth and processing resources of the routers involved.

The following diagram illustrates how RRs can improve the situation when only six routers are involved. The AS without RRs requires 15 sessions between the routers. In the AS with RRs, the two RRs receive route updates from the reflector clients (unlabeled routers in the diagram) in their cluster, as well as other RRs, and pass them on to the border router. The RR configuration requires only six sessions. This example shows a reduction of 60% for the number of required sessions.

Required sessions within an AS with and without RRs



The BGP commands related to RRs include:

```
config router bgp
  config neighbor
    set route-reflector-client {enable | disable}
    set route-server-client {enable | disable}
  end
end
```

Confederations

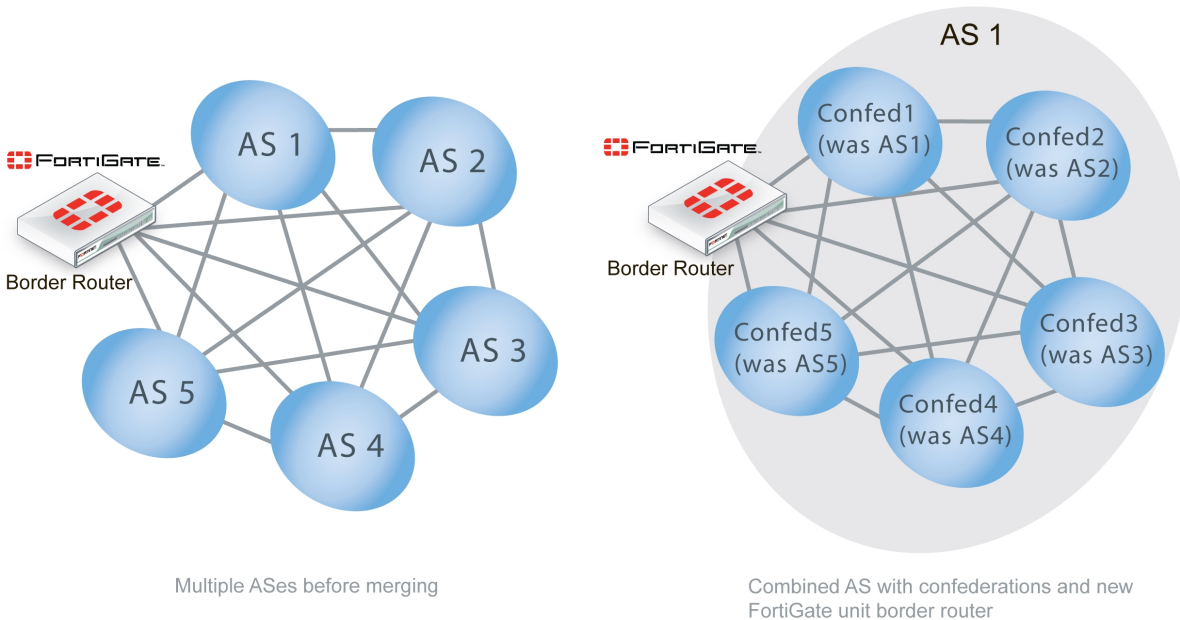
Confederations were introduced to reduce the number of BGP advertisements on a segment of the network and reduce the size of the routing tables. Confederations essentially break up an AS into smaller units.

Confederations are defined in [RFC 3065](#) and [RFC 1965](#).

Within a confederation, all routers communicate with each other in a full mesh arrangement. Communications between confederations is more like inter-AS communications because many of the attributes are changed as they would be for BGP communications leaving the AS, or eBGP.

Confederations are useful when merging ASs. Each AS being merged can easily become a confederation, which requires few changes. Any additional permanent changes can then be implemented over time, as required. The diagram below shows the group of ASs before merging and the corresponding confederations afterward, as part of the single AS with the addition of a new border router. It should be noted that after merging, if the border router becomes a route reflector, then each confederation only needs to communicate with one other router instead of five others.

AS merging using confederations



Confederations and RRs perform similar functions: they both sub-divide large ASs for more efficient operation. They differ in that route reflector clusters can include routers that aren't members of a cluster, whereas routers in a confederation must belong to that confederation. Also, confederations place their confederation numbers in the AS_PATH attribute, making it easier to trace.

It's important to note that while confederations essentially create sub-ASs, all the confederations within an AS appear as a single AS to external ASs.

Confederation related BGP commands include the following:

```
config router bgp
  set confederation-identifier <peerid_integer>
end
```

BGP conditional advertisements

Normally, routes are propagated regardless of the existence of a different path. The BGP conditional advertisement feature allows a route not to be advertised, based on the existence or non-existence of other routes. With this feature, a child table under `bgp.neighbor` is introduced. Any route matched by one of the route-maps specified in the table will be advertised to the peer, based on the corresponding route-map condition.

You can enable and disable conditional advertisements using the CLI.

To configure BGP conditional advertisements - CLI:

```
config router bgp
  set as 3
  config neighbor
    edit "10.10.10.10"
      set remote-as 3
      config conditional-advertise
        edit "route-map-to-match-sending"
```

```

        set condition-routemap "route-map-to-match-condition"
        set condition-type [exist | non-exist]
    next
end
next
end

```

BGP neighbor groups

The BGP neighbor group feature allows a large number of neighbors to be configured automatically based on a range of neighbors' source addresses.

To configure BGP neighbor groups - CLI:

Start by adding a BGP neighbor group:

```

config router bgp
  config neighbor-group
    edit <neighbor-group-name>
      set remote-as 100
    ...

```

(All options for BGP neighbor group are supported except `password`.)

```

end

```

Then add a BGP neighbor range:

```

config router bgp
  config neighbor-range
    edit 1
      set prefix 192.168.1.0/24
      set max-neighbor-num 100
      set neighbor-group <neighbor-group-name>
    next
  end

```

Network Layer Reachability Information

Network Layer Reachability Information (NLRI) is unique to BGP-4. It's sent as part of the update messages sent between BGP routers and contains information necessary to supernet, or aggregate route, information. The NLRI includes the length and prefix that, when combined, are the address of the aggregated routes referred to.

There is only one NLRI entry per BGP update message.

BGP attributes

Each route in a BGP network has a set of attributes associated with it. These attributes define the route and are modified, as required, along the route.

BGP can work well with mostly default settings, but if you're going to change settings you need to understand the roles of each attribute and how they affect those settings.

The BGP attributes include:

Attribute	Description
AS_PATH	A list of ASs a route has passed through. For more information, see "AS_PATH" on page 2194 .
MULTI_EXIT_DESC (MED)	Which router to use to exit an AS with more than one external connection. For more information, see "MULTI_EXIT_DESC" on page 2195 .
COMMUNITY	Used to apply attributes to a group of routes. For more information, see "COMMUNITY" on page 2195 .
NEXT_HOP	Where the IP packets should be forwarded to, like a gateway in static routing. For more information, see "NEXT_HOP" on page 2196 .
ATOMIC_AGGREGATE	Used when routes have been summarized to tell downstream routers not to de-aggregate the route. For more information, see "ATOMIC_AGGREGATE" on page 2196 .
ORIGIN	Used to determine if the route is from the local AS or not. For more information, see "ORIGIN" on page 2196 .
LOCAL_PREF	Used only within an AS to select the best route to a location (like MED).



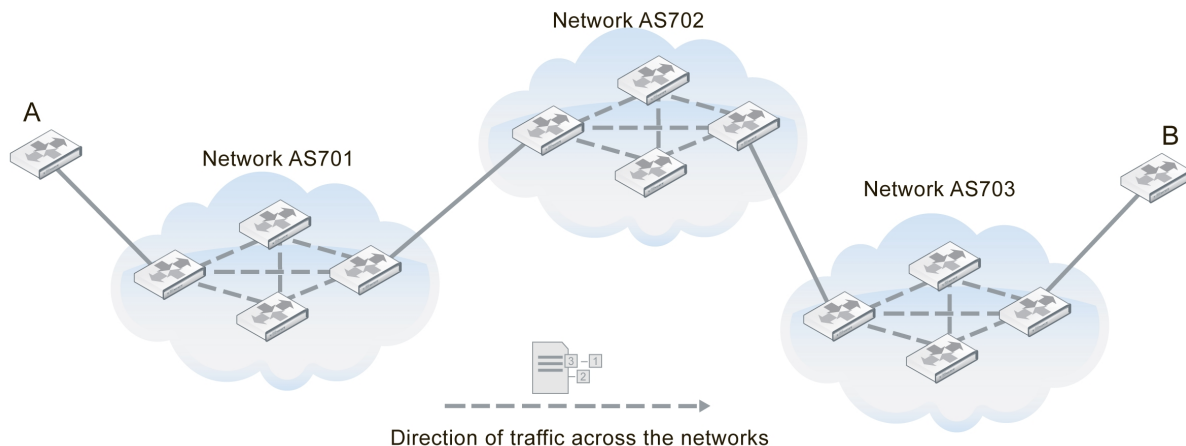
Inbound policies on FortiGate devices can change the NEXT-HOP, LOCAL-PREF, MED and AS-PATH attributes of an internal BGP (iBGP) route for its local route selection purposes. However, outbound policies on the device can't affect these attributes.

AS_PATH

AS_PATH is the BGP attribute that keeps track of each AS that a route advertisement has passed through. AS_PATH is used by confederations and by exterior BGP (EBGP) to help prevent routing loops. A router knows there is a loop if it receives an AS_PATH with that router's AS in it. The diagram below shows the route between Router A and Router B. The AS_PATH from A to B would read 701,702,703 for each AS that the route passes through.

As of the beginning of 2010, the industry upgraded from 2-byte to 4-byte AS_PATHs. This upgrade was due to the imminent exhaustion of 2-byte AS_PATH numbers. FortiOS supports 4-byte AS_PATHs in its BGP implementation.

AS_PATH of 701,702, 703 between routers A and B



The BGP commands related to AS_PATH include the following:

```
config router bgp
  set bestpath-as-path-ignore {enable | disable}
end
```

MULTI_EXIT_DESC

BGP AS systems can have one or more routers that connect them to other ASs. For ASs with more than one connecting router, the Multi-Exit Discriminator (MED) lists which router is best to use when leaving the AS. The MED is based on attributes, such as delay. It's a recommendation only, as some networks may have different priorities.

BGP updates advertise the best path to a destination network. When a FortiGate receives a BGP update, the FortiGate examines the MED attribute of potential routes to determine the best path to a destination network before recording the path in the local FortiGate routing table.

FortiGate devices have the option to treat any routes without an MED attribute as the worst possible routing choice. This can be useful because a lack of MED information is a lack of routing information, which can be suspicious as a possible hacking attempt or an attack on the network. At best, it signifies an unreliable route to select.

The BGP commands related to MED include the following:

```
config router bgp
  set always-compare-med {enable | disable}
  set bestpath-med-confed {enable | disable}
  set bestpath-med-missing-as-worst {enable | disable}
  set deterministic-med {enable | disable}
  config neighbor
    set attribute-unchanged [as-path] [med] [next-hop]
  end
end
```

COMMUNITY

A community is a group of routes that have the same routing policies applied to them. This saves time and resources. A community is defined by the COMMUNITY attribute of a BGP route.

A FortiGate can set the COMMUNITY attribute of a route to assign the route to predefined paths (see [RFC 1997](#)). The FortiGate can examine the COMMUNITY attribute of learned routes to perform local filtering and/or redistribution.

The BGP commands related to COMMUNITY include the following:

```
config router bgp
    set send-community {both | disable | extended | standard}
end
```

NEXT_HOP

The NEXT_HOP attribute says what IP address the packets should be forwarded to next. Each time the route is advertised, this value is updated. The NEXT_HOP attribute is much like a gateway in static routing.

FortiGate devices allow you to change the advertising of the FortiGate device's IP address (instead of the neighbor's IP address) in the NEXT_HOP information that is sent to IBGP peers. This is changed with the `config neighbor, set next-hop-self` command.

The BGP commands related to NEXT_HOP include the following:

```
config router bgp
    config neighbor
        set attribute-unchanged [as-path] [med] [next-hop]
        set next-hop-self {enable | disable}
    end
end
```

ATOMIC_AGGREGATE

The ATOMIC_AGGREGATE attribute is used when routes have been summarized. It indicates which AS and which router summarize the routes. It also tells downstream routers not to de-aggregate the route. Summarized routes are routes with similar information that have been combined, or aggregated, into one route that's easier to send in updates for. When it reaches its destination, the summarized routes are split back up into the individual routes.

The FortiGate doesn't specifically set this attribute in the BGP router command, but it's used in the route map command.

The CLI commands related to ATOMIC_AGGREGATE include the following:

```
config router route-map
    edit <route_map_name>
        config rule
            edit <route_map_rule_id>
                set set-aggregator-as <id_integer>
                set set-aggregator-ip <address_ipv4>
                set set-atomic-aggregate {enable | disable}
            end
        end
    end
```

ORIGIN

The ORIGIN attribute records where the route came from. The options can be IBGP, EBGP, or incomplete. This information is important because internal routes (IBGP) are, by default, higher priority than external routes (EBGP). However, incomplete ORIGINS are the lowest priority of the three.

The CLI commands related to ORIGIN include the following:

```
config router route-map
  edit <route_map_name>
    set comments <string>
    config rule
      edit <route_map_rule_id>
        set match-origin {egp | igp | incomplete | none}
      end
    end
  end
end
```

How BGP works

BGP is a link-state routing protocol and keeps link-state information about the status of each network link it has connected. A BGP router receives information from its peer routers that have been defined as neighbors. BGP routers listen for updates from these configured neighboring routers on TCP port 179.

A BGP router is a finite state machine with six various states for each connection. As two BGP routers discover each other and establish a connection, they go from the idle state and through the various states until they reach the established state. An error can cause the connection to drop and the state of the router to reset to either active or idle. These errors can be caused by TCP port 179 not being open, a random TCP port above port 1023 not being open, the peer address being incorrect, or the AS number being incorrect.

When BGP routers start a connection, they negotiate which (if any) optional features will be used, such as multiprotocol extensions, that can include IPv6 and VPNs.

IBGP versus EBGp

When you read about BGP, you often see EBGp or IBGP mentioned. These are both BGP routing, but BGP used in different roles. Exterior BGP (EBGP) involves packets crossing multiple autonomous systems (ASs) and interior BGP (IBGP) involves packets that stay within a single AS. For example, the AS_PATH attribute is only useful for EBGp where routes pass through multiple ASs.

These two modes are important because some features of BGP are used only for one of EBGp or IBGP. For example, confederations are used in EBGp and RRs are used only in IBGP. Also, routes learned from IBGP have priority over routes learned from EBGp.

FortiGate devices have some commands that are specific to EBGp, including the following:

- automatically resetting the session information to external peers if the connection goes down: `set fast-external-failover {enable | disable}`
- setting an administrative distance for all routes learned from external peers (you must also configure local and internal distances if this is set): `set distance-external <distance_integer>`
- enforcing EBGp multihops and their TTL (number of hops): `set ebgp-enforce-multihop {enable | disable}` and `set ebgp-multihop-ttl <seconds_integer>`

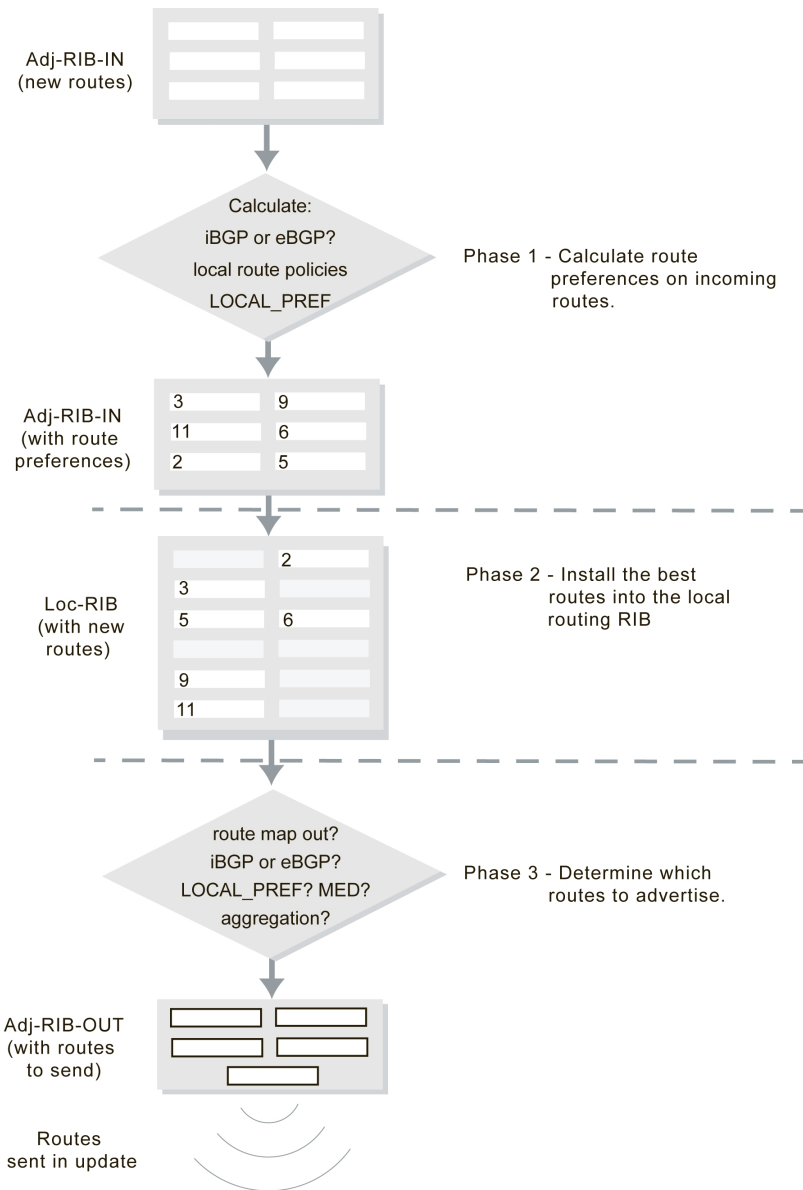
BGP path determination: which route to use

Firstly, recall that the number of available or supported routes isn't set by the configuration but depends on the available memory on the FortiGate. All learned routes and their attributes come into the BGP router in raw form. Before routes are installed in the routing table or are advertised to other routers, three levels of decisions must be made.

The three phases of BGP best path determination don't change. However, some manufacturers have added more information to the process, such as Cisco's WEIGHT attribute, to allow an administrator to force one route's selection over another.

There is one Adj-RIB-IN and Adj-RIB-OUT for each configured neighbor. They are updated when the FortiGate receives BGP updates or when the FortiGate sends out BGP updates.

The three phases of a BGP routing decision



Decision phase 1

At this phase, the decision is to calculate how preferred each route and its NRI are the Adjacent Routing Information Base Incoming (Adj-RIBs-In) compared to the other routes. For internal routes (IBGP), policy information or LOCAL_PREF is used. For external peer learned routes, it's based strictly on policy. These rules set up a list of which routes are most preferred going into Phase 2.

Decision phase 2

Phase 2 involves installing the best route to each destination into the local Routing Information Base (Loc-RIB). Effectively, the Loc-RIB is the master routing table. Each route from Phase 1 has their NEXT_HOP checked to ensure the destination is reachable. If it's reachable, the AS_PATH is checked for loops. After that, routes are installed based on the following decision process:

- If there's only one route to a location, it's installed.
- If there are multiple routes to the same location, use the most preferred route from Level 1.
- If there's a tie, break the tie based on the following, in descending order of importance: shortest AS_PATH, smallest ORIGIN number, smallest MED, EBGP over IBGP, smallest metric or cost for reaching the NEXT_HOP, BGP identifier, and lowest IP address.

Note that the new routes that are installed into the Loc-RIB are in addition to any existing routes in the table. Once Phase 2 is completed, the Loc-RIB will consist of the best of both the new and older routes.

Decision phase 3

Phase 3 is route distribution or dissemination. This is the process of deciding which routes the router will advertise. If there's any route aggregation or summarizing, it happens here. Also, any route filtering from route maps happens here.

Once Phase 3 is complete, an update can be sent out to update the neighbor of new routes.

Aggregate routes and addresses

BGP-4 allows classless routing, which uses netmasks as well as IP addresses. This classless routing allows the configuration of aggregate routes by stating the address bits the aggregated addresses have in common. For more information, see ["BGP" on page 1](#).

The ATOMIC_AGGREGATE attribute informs routers that the route has been aggregated and shouldn't be de-aggregated. An associated AGGREGATOR attribute include the information about the router that did the aggregating including its AS.

The BGP commands associated with aggregate routes and addresses are the following:

```
config router bgp
  config aggregate-address
    edit <aggr_addr_id>
      set as-set {enable | disable}
      set prefix <address_ipv4mask>
      set summary-only {enable | disable}
    end
  config aggregate-address6
    edit <aggr_addr_id>
      set as-set {enable | disable}
      set prefix6 <address_ipv6mask>
      set summary-only {enable | disable}
    end
```

Configuring BGP graceful restart process on timer

You can configure the BGP graceful restart process to stop only when the restart timer expires, using the following CLI commands:

```
config router bgp
```

```
set graceful-end-on-timer enable
```

Configuring option to bring down BGP neighbor when the link is down

You can configure an option to bring down BGP neighbors when the outgoing interface is down, using the following CLI commands:

```
config router bgp
config neighbor
edit <ip_address>
set linkdown-failover enable
```

Configuring option to keep routes for a period after the BGP neighbor is down

You can configure an option to keep routes for a period after the BGP neighbor is down. If you enable this option for a BGP neighbor, the route learned from the neighbor is kept for the configured `graceful-stalepath-time` after the neighbor is down because of hold timer expiration or TCP connection failure.

To configure this option, use the following CLI commands:

```
config router bgp
config neighbor
edit <ip_address>
set stale-route enable
```

BGP local-AS support

A FortiGate supports BGP local-AS. Local-AS allows you to configure a BGP speaker to have a real local-AS and a secondary local-AS for a specific neighbor, so its local-AS number appears different to different neighbors.

You can configure a BGP speaker to have a real local-AS and a secondary local-AS for a specific neighbor, so the local-AS number appears different to neighbor B and neighbor A.

To configure BGP local-AS for BGP peers - CLI:

```
config router bgp
config neighbor
edit "neighbor" / edit <ip_address>
...
set local-as 300 (?) / set local-as <integer>
set local-as-no-prepend {enable | disable}
set local-as-replace-as {enable | disable}
end
```

CLI option	Description
<ip_address>	The IP/IPv6 address of neighbor
<local-as <integer>>	The local-AS number
local-as-no-prepend { enable disable }	Set this to enable if you do not want to prepend local-AS to incoming updates.
local-as-replace-as { enable disable }	Set this to enable to replace a real AS with local-AS in outgoing updates.

Troubleshooting BGP

There are some features in BGP that are used to deal with problems that may arise. Typically, the problems with a BGP network that has been configured involve routes going offline frequently. This is called route flap and causes problems for the routers using that route.

Clearing routing table entries

To see if a new route is being properly added to the routing table, you can clear all or some BGP neighbor connections (sessions) using the `execute router clear bgp` command.

For example, if you have 10 routes in the BGP routing table and you want to clear the specific route to IP address 10.10.10.1, enter the following CLI command:

```
execute router clear bgp ip 10.10.10.1
```

To remove all routes for AS number 650001, enter the following CLI command:

```
execute router clear bgp as 650001
```

Route flap

When routers or hardware along a route go offline and back online that is called a route flap. Flapping is the term that is used if these outages continue, especially if they occur frequently.

Route flap is a problem in BGP because each time a peer or a route goes down, all the peer routers that are connected to that out-of-service router advertise the change in their routing tables. This creates a lot of administration traffic on the network and the same traffic re-occurs when that router comes back online. If the problem is something like a faulty network cable that wobbles online and offline every 10 seconds, there could easily be an overwhelming amount of routing updates sent out unnecessarily.

Another possible reason for route flap occurs with multiple FortiGate devices in HA mode. When an HA cluster fails over to the secondary unit, other routers on the network may see the HA cluster as being offline, resulting in route flap. While this doesn't occur often, or more than once at a time, it can still result in an interruption in traffic which is unpleasant for network users. The easy solution for this problem is to increase the timers on the HA cluster, such as TTL timers, so they don't expire during the failover process. Also, configuring graceful restart on the HA cluster helps with a smooth failover.

The first method of dealing with route flap is to check your hardware. If a cable is loose or bad, it can easily be replaced and eliminate the problem. If an interface on the router is bad, either avoid using that interface or swap in a functioning router. If the power source is bad on a router, either replace the power supply or use a power conditioning backup power supply. These quick and easy fixes can save you from configuring more complex BGP options. However, if the route flap is from another source, configuring BGP to deal with the outages will ensure your network users uninterrupted service.

Some methods of dealing with route flap in BGP include:

- [Hold down timer](#)
- [Dampening](#)
- [Graceful restart](#)
- [Troubleshooting BGP](#)

Hold down timer

The first line of defense to a flapping route is the hold down timer. This timer reduces how frequently a route going down will cause a routing update to be broadcast.

Once activated, the hold down timer won't allow the FortiGate to accept any changes to that route for the duration of the timer. If the route flaps five times during the timer period, only the first outage will be recognized by the FortiGate. For the duration of the other outages, there won't be changes because the Fortigate is essentially treating this router as down. If the route is still flapping after the timer expires, it'll happen all over again.

Even if the route isn't flapping (for example, if it goes down, comes up, and stays back up) the timer still counts down and the route is ignored for the duration of the timer. In this situation, the route is seen as down longer than it really is but there will be only the one set of route updates. This isn't a problem in normal operation because updates aren't frequent.

Also, the potential for a route to be treated as down when it's really up can be viewed as a robustness feature. Typically, you don't want most of your traffic being routed over an unreliable route. So if there's route flap going on, it's best to avoid that route if you can. This is enforced by the hold down timer.

How to configure the hold down timer

There are three different route flapping situations that can occur: the route goes up and down frequently, the route goes down and back up once over a long period of time, or the route goes down and stays down for a long period of time. These can all be handled using the hold down timer.

For example, your network has two routes that you want to set the hold down timer for. One is your main route (to 10.12.101.4) that all of your Internet traffic goes through, and it can't be down for long if it's down. The second is a low speed connection to a custom network that's used infrequently (to 10.13.101.4). The hold down timer for the main route should be fairly short, let's say 60 seconds instead of the default 180 seconds. The second route timer can be left at the default, or even longer since it's rarely used. In your BGP configuration this looks like the following:

```
config router bgp
  config neighbor
    edit 10.12.101.4
      set holddown-timer 60
    next
    edit 10.13.101.4
      set holddown-timer 180
    next
  end
end
```

Dampening

Dampening is a method that's used to limit the amount of network problems due to flapping routes. With dampening, the flapping still occurs but the peer routers pay less and less attention to that route as it flaps more often. One flap doesn't start dampening, but the second flap starts a timer where the router won't use that route because it's considered unstable. If the route flaps again before the timer expires, the timer continues to increase. There's a period of time called the reachability half-life, after which a route flap will be suppressed for only half the time. This half-life comes into effect when a route has been stable for a while but not long enough to clear all the dampening completely. For the flapping route to be included in the routing table again, the suppression time must expire.

If the route flapping was temporary, you can clear the flapping or dampening from the FortiGate device's cache by using one of the `execute router clear bgp` CLI commands:

```
execute router clear bgp dampening {<ip_address> | <ip/netmask>}
or
execute router clear bgp flap-statistics {<ip> | <ip/netmask>}
```

For example, to remove route flap dampening information for the 10.10.0.0/16 subnet, enter the following CLI command:

```
execute router clear bgp dampening 10.10.0.0/16
```

The BGP commands related to route dampening are the following:

```
config router bgp
  set dampening {enable | disable}
  set dampening-max-suppress-time <minutes_integer>
  set dampening-reachability-half-life <minutes_integer>
  set dampening-reuse <reuse_integer>
  set dampening-route-map <routemap-name_str>
  set dampening-suppress <limit_integer>
  set dampening-unreachability-half-life <minutes_integer>
end
```

Graceful restart

BGP4 has the capability to gracefully restart.

In some situations, route flap is caused by routers that appear to be offline but the hardware portion of the router (control plane) can continue to function normally. One example of this is when some software is restarting or being upgraded but the hardware can still function normally.

Graceful restart is best used for these situations where routing won't be interrupted, but the router is unresponsive to routing update advertisements. Graceful restart doesn't have to be supported by all routers in a network, but the network will benefit when more routers support it.



FortiGate HA clusters can benefit from graceful restart. When a failover takes place, the HA cluster advertises that it's going offline, and won't appear as a route flap. It will also enable the new HA main unit to come online with an updated and usable routing table. If there's a flap, the HA cluster routing table will be out-of-date.

For example, the FortiGate is one of four BGP routers that send updates to each other. Any of those routers may support graceful starting. When a router plans to go offline, it sends a message to its neighbours stating how long it expects to be offline. This way, its neighboring routers don't remove it from their routing tables. However, if that router isn't back online when expected, the routers will mark it offline. This prevents routing flap and its associated problems.

Scheduled time offline

Graceful restart is a means for a router to advertise that it is going to have a scheduled shutdown for a very short period of time. When neighboring routers receive this notice, they will not remove that router from their routing table until after a set time elapses. During that time, if the router comes back online, everything continues to function as normal. If that router remains offline longer than expected, then the neighboring routers will update their routing tables as they assume that router will be offline for a long time.

FortiGate devices support both graceful restart of their own BGP routing software and neighboring BGP routers.

For example, if a neighbor of the FortiGate with an IP address of 172.20.120.120 supports graceful restart, enter the following CLI command:

```
config router bgp
  config neighbor
    edit 172.20.120.120
      set capability-graceful-restart enable
    end
  end
```

If you want to configure graceful restart on the FortiGate where you expect the Fortigate to be offline for no more than two minutes, and after three minutes the BGP network should consider the FortiGate to be offline, enter the following CLI commands:

```
config router bgp
  set graceful-restart enable
  set graceful-restart-time 120
  set graceful-stalepath-time 180
end
```

The BGP commands related to BGP graceful restart are the following:

```
config router bgp
  set graceful-restart { disable | enable }
  set graceful-restart-time <seconds_integer>
  set graceful-stalepath-time <seconds_integer>
  set graceful-update-delay <seconds_integer>
  config neighbor
    set capability-graceful-restart { enable | disable }
  end
end

execute router restart
```

Before the restart, the router sends its peers a message to say it's restarting. The peers mark all the restarting router's routes as stale, but they continue to use the routes. The peers assume the router will restart, check its routes, and take care of them, if needed, after the restart is complete. The peers also know what services the restarting router can maintain during its restart. After the router completes the restart, the router sends its peers a message to say it's done restarting.

BFD

Bidirectional Forwarding Detection (BFD) is a protocol that you can use to quickly locate hardware failures in the network. Routers running BFD communicate with each other and if a timer runs out on a connection then that router is declared down. BFD then communicates this information to the routing protocol and the routing information is updated. For more information about BFD, see ["BFD" on page 2106](#).

Dual-homed BGP example

This is an example of a small network that uses BGP routing connections to two ISPs. This is a common configuration for companies that need redundant connections to the Internet for their business.

This configuration is for a small company connected to two ISPs. The company has one main office, the Head Office, and uses static routing for internal routing on that network.

Both ISPs use BGP routing and connect to the Internet directly. They want the company to connect to the ISP networks using BGP. They also use graceful restart to prevent updates that aren't needed and use smaller timer values to detect network failures faster.

As can be expected, the company wants to keep their BGP configuration relatively simple and easy to manage. The current configuration has only 3 routers to worry about: the 2 ISP border routers and the FortiGate. This means that the FortiGate has only two neighbor routers to configure.

This configuration has the added benefit of being easy to expand if the company wants to add a remote office in the future.

To keep the configuration simple, the company is allowing only HTTP, HTTPS, FTP, and DNS traffic out of the local network. This allows employees access to the Internet and their web mail.

Why dual home?

Dual homing means having two separate independent connections to the Internet. Servers in this configuration have also been called bastion hosts and can include DNS servers which require multiple connections.

Benefits of dual homing can include:

- Redundant Internet connection that essentially never fails
- Faster connections through one ISP or the other for some destinations, such as other clients of those ISPs
- Load balancing traffic to the company network
- Easier to enable more traffic through two connections than upgrading one connection to bigger bandwidth
- Easier to create protection policies for different traffic through a specific ISP

Some companies require reliable Internet access at all times as part of their business. Consider a doctor operating remotely who has their Internet connection fail — the consequences can easily be life or death.

Dual homing is an extra expense for the second ISP connection and more work to configure and maintain the more complex network topology.

Potential dual homing issues

BGP comes with load balancing issues and dual homing is in the same category. BGP doesn't inherently deal well with load balancing or getting default routes through BGP. Ideally, one connection may be best for certain destinations but it may not have that traffic routed to it, which makes the load balancing less than perfect. This kind of fine tuning can be very time consuming and usually results in a best effort situation.

When dual homing isn't configured properly, your network may become a link between your ISPs and result in very high traffic between the ISPs that doesn't originate from your network. The problem with this situation is that your traffic may not have the bandwidth it needs, and you'll also be paying for a large volume of traffic that isn't yours. This problem can be solved by not broadcasting or redistributing BGP routes between the ISPs.

If you learn your default routes from the ISPs, in this example, you may run into an asymmetric routing problem where your traffic loops out one ISP and back to you through the other ISP. If you think this may be happening, you can turn on asymmetric routing on the FortiGate (`config system settings, set asymmetric enable`) to verify if that's the problem. Turn this feature off once this is established, since it disables many features on the FortiGate by disabling stateful inspection. Solutions to this problem can include using static routes for default routes instead of learning them through BGP or configuring VDOMs on the FortiGate to provide a slightly different path back that isn't a true loop.


Network layout and assumptions

The network layout for the basic BGP example involves the company network being connected to both ISPs as shown below. In this configuration, the FortiGate is the BGP border router between the Company AS, ISP1's AS, and ISP2's AS.

The components of the layout include the following:

- The Company AS (AS number 1) is connected to ISP1 and ISP2 through the FortiGate.
- The Company has one internal network: the Head Office network at 10.11.101.0/24.
- The FortiGate internal interface is on the company's internal network with an IP address of 10.11.101.110.
- The FortiGate external1 interface is connected to ISP1's network with an IP address of 172.20.111.5, which is an address supplied by the ISP.
- The FortiGate external2 interface is connected to IPS2's network with an IP address of 172.20.222.5, which is an address supplied by the ISP.
- ISP1 AS has an AS number of 650001 and ISP2 has an AS number of 650002.
- Both ISPs are connected to the Internet.
- The ISP1 border router is a neighbor (peer) of the FortiGate. It has an address of 172.21.111.4.
- The ISP2 border router is a neighbor (peer) of the FortiGate. It has an address of 172.22.222.4.
- Apart from graceful restart and shorter timers (holdtimer and keepalive), default settings are to be used whenever possible.

Basic BGP network topology



Assumptions

The basic BGP configuration procedure follows these assumptions:

- ISP1 is the preferred route and ISP2 is the secondary route
- All basic configuration can be completed in both the GUI and CLI
- Only one AS is used for the company

For these reasons, this example configuration does not include:

- Bidirectional forwarding detection
- Route maps
- Access lists
- Changing redistribution defaults (make link when example is set up)
- IPv6

For more information about these features, see the corresponding section.

Configuring the FortiGate

In this topology, the FortiGate is the link between the company network and the ISP network. The FortiGate is the only BGP router on the company network, but there's at least one other BGP router on the ISP network. There may be more BGP routers, but we don't have that information.

As mentioned in the general configuration steps, the ISP must be notified of the company's BGP router configuration when complete as it will need to add the FortiGate BGP router as a neighbor router on its domain. This step is required for the FortiGate to receive BGP routing updates from the ISP network and outside networks.

If the ISP has any special BGP features enabled, such as graceful restart or route dampening, that should be determined ahead of time so those features can be enabled on the FortiGate.

To configure the FortiGate as a BGP router:

1. [Configure interfaces and default routes](#)
2. [Configure firewall services, addresses, and policies](#)
3. [Set the FortiGate BGP information](#)
4. [Add the internal network to the AS](#)
5. [Additional FortiGate BGP configuration](#)

Configure interfaces and default routes

The FortiGate is connected to three networks: the company network on the internal interface, the ISP1 network on the external1 interface, and the ISP2 network on the external2 interface.

This example uses basic interface settings. Check with your ISP to determine if additional settings are required, such as setting the maximum MTU size or if gateway detection is supported.

High end FortiGate models don't have interfaces labeled as Internal or External. Instead, for clarity, we're using the alias feature to name interfaces for these roles.

Default routes to both external interfaces are configured here also. Both are needed in case one goes offline. ISP1 is the primary connection and has a smaller administrative distance so it will be preferred over ISP2. Both distances are set low so they will be preferred over any learned routes.

To configure the FortiGate interfaces - GUI:

1. Go to **Network > Interfaces**.
2. Edit port 1 (internal) interface.
3. Set the following information and select **OK**.

Alias	internal
IP/Network Mask	10.11.101.110/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	Company internal network
Interface State	Enabled

4. Edit port 2 (external1) interface.
5. Set the following information and select **OK**.

Alias	external1
IP/Network Mask	172.21.111.5/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	ISP1 External BGP network
Interface State	Enabled

6. Edit port 3 (external2) interface.
7. Set the following information and select **OK**.

Alias	external2
IP/Network Mask	172.22.222.5/255.255.255.0
Administrative Access	HTTPS SSH PING
Comments	ISP2 External BGP network
Interface State	Enabled

To configure the FortiGate interfaces - CLI:

```
config system interface
  edit port1
    set alias internal
```

```

        set ip 10.11.101.110 255.255.255.0
        set allowaccess http https ssh
        set description "Company internal network"
        set status up
    next
    edit port2
        set alias external1
        set ip 172.21.111.5 255.255.255.0
        set allowaccess https ssh
        set description "ISP1 External BGP network"
        set status up
    next
    edit port3
        set alias external2
        set ip 172.22.222.5 255.255.255.0
        set allowaccess https ssh
        set description "ISP2 External BGP network"
        set status up
    next
end

```

To configure default routes for both ISPs - GUI:

1. Go to **Network > Static Routes**.
2. Delete any existing routes with a IP/Mask of address of 0.0.0.0/0.0.0.0
3. Select **Create New** and set the following information.

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.21.111.5
Interface	port2
Administrative Distance	10

4. Select **OK**.
5. Select **Create New** and set the following information.

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	172.22.222.5
Interface	port3
Administrative Distance	15

6. Select **OK**.

To configure default routes for both ISPs - CLI:

```

config router static
edit 1
    set device "port2"
    set distance 10

```



```
        set gateway 172.21.111.5
    next
    edit 2
        set device "port3"
        set distance 15
        set gateway 172.22.222.5
    next
end
```

Configure firewall services, addresses, and policies

To create the security policies, you create the firewall services group that will include all the services that will be allowed, define the addresses that will be used in the security policies, and configure the security policies themselves.

To keep the configuration simple, the company is allowing only HTTP traffic out of the local network. This will allow employees access to the Internet and their web mail. DNS services will also be allowed through the firewall.

The security policies will allow HTTP traffic (port 80 and port 8080), HTTPS traffic (port 443), FTP traffic (port 21), and DNS traffic (port 53 and port 953) in both directions. Also, BGP (port 179) may need access through the firewall.



For added security, you may want to define a smaller range of addresses for the internal network. For example, if only 20 addresses are used, only allow those addresses in the range.

To keep things simple, a zone is used to group the two ISP interfaces together. This allows for the use of one security policy to apply to both ISPs at the same time. Remember to block intra-zone traffic as this helps to prevent one ISP sending traffic to the other ISP through the FortiGate, using your bandwidth. The zone keeps configuration simple and if there's a need for separate policies for each ISP in the future, they can be created and the zone can be deleted.

The addresses that will be used are the addresses of the FortiGate internal and external ports and the internal network.

More policies or services can be added in the future as applications are added to the network. For more information about security policies, see the [FortiOS Firewall Handbook](#).



When configuring security policies, always enable logging to help you track and debug your traffic flow.

To create a firewall services group - GUI:

1. Go to **Policy & Objects > Services**, select the dropdown arrow next to **Create New** and select **Service Group**.
2. For **Group Name**, enter "Basic_Services".
3. From the **Members** dropdown, choose the following six services: BGP, FTP, FTP_GET, FTP_PUT, DNS, HTTP, and HTTPS.
4. Select **OK**.

To create a firewall services group - CLI:

```
config firewall service group
```

```

edit "Basic_Services"
    set member "BGP" "DNS" "FTP" "FTP_GET" "FTP_PUT" "HTTP" "HTTPS"
next
end

```

To create a zone for the ISP interfaces - GUI:

1. Go to **Network > Interfaces**.
2. Select the caret to the right of **Create New** and then select **Zone**.
3. Enter the following information:

Name	ISPs
Block intra-zone traffic	enable
Interface Members	port2 port3

4. Select **OK**.

To create a zone for the ISP interfaces - CLI:

```

config system zone
    edit "ISPs"
        set interface "port2" "port3"
        set intrazone block
    next
end

```

To add the firewall addresses - GUI:

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New** and set the following information:

Category	Address
Name	Internal_network
Type	Subnet / IP Range
Subnet / IP Range	10.11.101.0 255.255.255.0
Interface	port1

3. Select **OK**.

To add the firewall addresses - CLI:

```

config firewall address
    edit Internal_network
        set associated-interface port1
        set subnet 10.11.101.0 255.255.255.0
    next
end

```

To add the HTTP and DNS security policies - GUI:

1. Go to **Policy & Objects > IPv4 Policy**, and select **Create New**.
2. Set the following information:

Incoming Interface	port1(internal)
Outgoing Interface	ISPs
Source	Internal_network
Destination	All
Schedule	Always
Service	Basic_services
Action	ACCEPT
Firewall / Network Options	Enable NAT
Log Allowed Traffic	Enable
Comments	ISP1 basic services out policy

3. Select **OK**.
4. Select **Create New** and set the following information:

Incoming Interface	ISPs
Outgoing Interface	port1(internal)
Source	All
Destination	Internal_network
Schedule	Always
Service	Basic_services
Action	ACCEPT
Firewall / Network Options	Enable NAT
Log Allowed Traffic	Enable
Comments	ISP1 basic services in policy

To add the security policies - CLI:

```
config firewall policy
edit 1
```

```

        set srcintf "port1"
        set srcaddr "Internal_network"
        set dstintf "ISPs"
        set dstaddr "all"
        set schedule "always"
        set service "Basic_services"
        set action accept
        set nat enable
        set profile-status enable
        set logtraffic enable
        set comments "ISP1 basic services out policy"
    next
edit 2
    set srcintf "ISPs"
    set srcaddr "all"
    set dstintf "port1"
    set dstaddr "Internal_network"
    set schedule "always"
    set service "Basic_services"
    set action accept
    set nat enable
    set profile-status enable
    set logtraffic enable
    set comments "ISP1 basic services in policy"
next
end

```

Set the FortiGate BGP information

When using the default information, there are only two fields to set to configure the FortiGate as a BGP router.

For this configuration, the FortiGate will be in a stub area with one route out — the ISP BGP router. Until you configure the ISP router as a neighbor, even that route out isn't available. So, while after this part of the configuration is complete, the FortiGate will be running BGP, it won't know about any other routers running BGP until the next part of the configuration is complete.

To set the BGP router information - GUI:

1. Go to **Network > BGP**.
2. Set the following information and select **OK**.

Local AS	1
Router ID	10.11.101.110

To set the BGP router information - CLI:

```

config router BGP
    set as 1
    set router-id 10.11.101.110
end

```

Add the internal network to the AS

The company is one AS with the FortiGate configured as the BGP border router connecting that AS to the two ISPs ASs. The internal network in the Company's AS must be defined. If there were other networks in the company, such as regional offices, they would be added here as well.

To set the networks in the AS - GUI:

1. Go to **Network > BGP**.
2. Under **Networks**, set the **IP/Netmask** to 10.11.101.0/255.255.255.0 and select **Add**.

To set the networks in the AS - CLI:

```
config router bgp
  config network
  edit 1
    set prefix 10.11.101.0 255.255.255.0
  next
end
end
```

Add BGP neighbor information

The configuration won't work unless you set **Remote AS** neighbors. This can be done in either the GUI or the CLI.

To configure the BGP neighbors - GUI:

1. Go to **Network > BGP**.
2. Add a **Neighbors IP** of 172.21.111.4 with the **Remote AS** set to 650001, then click **Add/Edit**.
3. Add another **Neighbors IP** of 172.22.222.4 with the **Remote AS** set to 650002, then click **Add/Edit**.

To configure the BGP neighbors - CLI:

```
config router BGP
  set as 1
  config neighbor
  edit "172.21.111.4"
    set remote-as 650001
  next
  edit "172.22.222.4"
    set remote-as 650002
  next
end
end
```

Additional FortiGate BGP configuration

At this point, those are all the settings that can be done in both the GUI and the CLI. The remaining configuration must be completed in the CLI.

These additional settings are mainly determined by your ISP requirements. They will determine your timers, such as keepalive timers, if extended features like BFD and graceful restart are being used, and so on. For this

example, some common simple features are being used to promote faster detections of network failures, which will result in better service for the company's internal network users.

The ISPs don't require authentication between peer routers.

These commands will enable or modify the following features on the FortiGate and, where possible, on neighboring routers as well:

- `bestpath-med-missing-as-worst`: Treats a route without an MED as the worst possible available route due to expected unreliability
- `fast-external-failover`: Immediately reset the session information associated with BGP external peers if the link used to reach them goes down
- `graceful-restart*`: Advertise reboots to neighbors so they don't see the router as offline, wait before declaring them offline, and how long to wait when they reboot before advertising updates. These commands apply to neighbors and are part of the BGP capabilities. This prevents unneeded routing updates.
- `holdtime-timer`: How long the router will wait for a keepalive message before declaring a router offline. A shorter time will find an offline router faster.
- `keepalive-timer`: How often the router sends out keepalive messages to neighbor routers to maintain those sessions.
- `log-neighbor-changes`: Log changes to the status of neighbor routers. This can be useful for troubleshooting from both internal and external networks.
- `connect-timer`: How long (in seconds) the FortiGate will try to reach this neighbor before declaring it offline.
- `weight`: Used to prefer routes from one neighbor over the other. In this example, ISP1 is the primary connection so it's weighted higher than ISP2.

To configure additional BGP options - CLI:

```
config router bgp
  set bestpath-med-missing-as-worst enable
  set fast-external-failover enable
  set graceful-restart enable
  set graceful-restart-time 120
  set graceful-stalepath-time 180
  set graceful-update-delay 180
  set holdtime-timer 120
  set keepalive-timer 45
  set log-neighbor-changes enable
config neighbor
  edit 172.21.111.4
    set connect-timer 60
    set description "ISP1"
    set holdtime-timer 120
    set keepalive-timer 45
    set weight 250
  next
  edit 172.22.222.4
    set connect-timer 60
    set description "ISP2"
    set holdtime-timer 120
    set keepalive-timer 45
    set weight 100
  next
end
end
```

Configuring other networking devices

There are two other networking devices that need to be configured: the BGP routers for both ISPs.

The ISPs' routers must add the FortiGate as a neighbor so route updates can be sent in both directions. Note that ISP1 isn't directly connected to ISP2, that we're aware of.

Inform both of your ISPs of the FortiGate device's BGP information. Once they have configured their router, you can test your BGP connection to the Internet.

They will require your FortiGate device's IP address of the connected interface, the route ID, and your company's AS number.

Testing this configuration

With the dual-homed BGP configuration in place, you should be able to send and receive traffic, send and receive routes, and not have any routing loops. Testing the networks will confirm that things are working as expected.

In general, for routing, you need to look at the routing table on different routers to see what routes are being installed. You also need to sniff packets to see how traffic is being routed in real-time. These two sources of information will normally tell you what you need to know.

Testing of this example's network configuration should be completed in the following parts:

- [Testing network connectivity](#)
- [Verifying the FortiGate device's routing tables](#)
- [Verifying traffic routing](#)
- [Verifying the dual-homed side of the configuration](#)

Testing network connectivity

A common first step in testing a new network topology is to test to see if you can reach the Internet and other locations as expected. If not, you may be prevented by cabling issues, software, or other issues.

The easiest way to test connections is to use ping, once you ensure that all of the FortiGate interfaces and ISP routers have ping support enabled. Also, ensure that the security policies allow ping through the firewall.

Connections to test, in this example, are the internal network to ISP1's router or the Internet, and the same for ISP2. If you can connect on the external side of the Fortinet, try to ping the internal network. These three tests should prove your basic network connections are working.



Once you've finished testing the network connectivity, turn off ping support on the external interfaces for additional security.

Verifying the FortiGate device's routing tables

The FortiGate routing table contains the routes that are stored for future use. If you're expecting certain routes to be there and they're not, this is a good indicator that your configuration isn't what you expected.

The `get router info routing-table details` CLI command will provide you with the routing protocol, destination address, gateway address, interface, and weighting for every route, as well as if the address is directly connected or not.

If you want to limit the display to BGP routes only, use the `get router info routing-table bgp` CLI command. If there are no BGP routes in the routing table, nothing will be displayed. In the CLI command, you can replace BGP with static, or other routing protocols, to only display those routes.

If you want to see the contents of the routing information database (RIB), use the `get router info routing-table database` CLI command. This will display the incoming routes that may or may not make it into the routing table.

Verifying traffic routing

Traffic may be reaching the internal network, but it may be using a different route than you think to get there.

Use a browser to try to access the Internet.

If needed, allow traceroute and other diag ports to be opened until things are working properly. Then remove access for them again.

Look for slow hops on the traceroute, or pings to a location, as they may indicate network loops that need to be fixed.

Any locations that have an unresolved traceroute or ping must be examined and fixed.

Use network packet sniffing to ensure traffic is being routed as you expect.

Verifying the dual-homed side of the configuration

Since there are two connections to the Internet in this example, theoretically you can pull the plug on one of the ISP connections, and all traffic will go through the other connection. Alternately, you may choose to remove a default route to one ISP, remove that ISP's neighbor settings, or change the weightings to prefer the other ISP. These alternate ways to test dual-homing don't change physical cabling, which may be preferred in some situations.

If this doesn't work as expected, things to check include:

- Default static routes: If these are wrong or don't exist, the traffic can't get out.
- BGP neighbor information: If the ISP router information is incorrect, the FortiGate won't be able to talk to it.

Redistributing and blocking routes in BGP

During normal BGP operation, peer routers redistribute routes from each other. However, in some specific situations it may be best not to advertise routes from one peer, such as if the peer is redundant with another peer (they share the same routes exactly), if it might be unreliable in some way or for some other reason. The FortiGate can also take routes it learns from other protocols and advertise them in BGP, for example OSPF or RIP. If your company hosts its own web or email servers, external locations will require routes to your networks to reach those services.

In this example, the company has an internal network in an OSPF area and is connected to a BGP AS and two BGP peers. The company goes through these two peers to reach the Internet. However, Peer 1 routes won't be advertised to Peer 2. The company internal user and server networks are running OSPF, and will redistribute those routes to BGP so external locations can reach the web and email servers.

Network layout and assumptions


The network layout for the BGP redistributing routes example involves the company network being connected to two BGP peers, as shown below. In this configuration, the FortiGate is the BGP border router between the

Company AS and the peer routers.

The components of the layout include:

- There's only one BGP AS in this example shared by the FortiGate and both peers: AS 65001.
- The company's FortiGate device connects to the Internet through two BGP peers.
- The company's internal networks on the dmz interface of the FortiGate with an IP of 10.11.201.0/24.
- The FortiGate device's interfaces are connected as follows:
 - port1 (dmz) has IP 10.11.201.110 and is the internal user and server network
 - port2 (external1) has IP 172.21.111.4 and is connected to Peer 1's network
 - port3 (external2) has IP 172.22.222.4 and is connected to Peer 2's network
- Peer 1 has IP 172.21.111.5, and Peer 2 has IP 172.22.222.5.
- OSPF Area 1 is configured on the dmz interface of the FortiGate, and is the routing protocol used by the internal users and servers.

BGP network topology



Assumptions

The BGP redistributing routes configuration procedure follows these assumptions:

- The FortiGate has been configured following the Install Guide
- Interfaces port1, port2, and port3 exist on the FortiGate
- We do not know the router manufacturers of Peer 1 and Peer 2
- We do not know what other devices are on the BGP AS or OSPF Area
- All basic configuration can be completed in both GUI and CLI

- Access lists and route maps will only be configured in CLI
- VDOMs are not enabled on the FortiGate

Configuring the FortiGate

1. [Configuring networks and firewalls on the FortiGate](#)
2. [Configuring BGP on the FortiGate](#)
3. [Configuring OSPF on the FortiGate](#)
4. [Configuring other networking devices](#)
5. [Configuring ECMP support for BGP](#)

Configuring networks and firewalls on the FortiGate

The FortiGate has three interfaces connected to networks: two external and one dmz.

Security policies must be in place to allow traffic to flow between these networks.

Firewall services will change depending on which routing protocol is being used on that network: either BGP or OSPF. Beyond that, all services that are allowed will be allowed in both directions due to the internal servers. The services allowed are web server services (DNS, HTTP, HTTPS, SSH, NTP, FTP*, SYSLOG, and MYSQL), email services (POP3, IMAP, and SMTP), and general troubleshooting services (PING, TRACEROUTE). To increase security, PING and TRACEROUTE can be removed once the network is up and working properly. Other services can be added later, as needed.

To configure the interfaces - GUI:

1. Go to **Network > Interfaces**.
2. Edit port1 (dmz) interface.
3. Set the following information and select **OK**.

Alias	dmz
IP/Network Mask	10.11.201.110/255.255.255.0
Administrative Access	HTTPS SSH PING
Description	OSPF internal networks
Administrative Status	Up

4. Edit port2 (external1) interface.
5. Set the following information and select **OK**.

Alias	external1
IP/Network Mask	172.21.111.4/255.255.255.0
Administrative Access	HTTPS SSH

Description	BGP external Peer 1
Administrative Status	Up

6. Edit port3 (external2) interface.
7. Set the following information and select **OK**.

Alias	external2
IP/Network Mask	172.22.222.4/255.255.255.0
Administrative Access	HTTPS SSH
Description	BGP external2 Peer2
Administrative Status	Up

To configure the FortiGate interfaces - CLI:

```
config system interface
  edit port1
    set alias dmz
    set ip 10.11.201.110 255.255.255.0
    set allowaccess https ssh ping
    set description "OSPF internal networks"
    set status up
  next
  edit port2
    set alias external1
    set ip 172.21.111.5 255.255.255.0
    set allowaccess https ssh
    set description "external1 Peer 1"
    set status up
  next
  edit port3
    set alias external2
    set ip 172.22.222.5 255.255.255.0
    set allowaccess https ssh
    set description "external2 Peer 2"
    set status up
  next
end
```

To configure the firewall addresses - GUI:

1. Go to **Policy & Objects > Objects > Addresses**.
2. Select **Create New** and set the following information.

Category	Address
Name	BGP_services

Type	Subnet / IP Range
Subnet / IP Range	10.11.201.0 255.255.255.0
Interface	port1

3. Select **OK**.

To configure the firewall addresses - CLI:

```
config firewall address
  edit "BGP_services"
    set associated-interface "port1"
    set subnet 10.11.201.0 255.255.255.0
  next
end
```

To configure firewall service groups - GUI:

1. Go to **Policy & Objects > Objects > Services**. Under the **Create New** dropdown menu, select **Service Group**.
2. Name the group BGP_Services.
3. Add the following services to the **Members** list: BGP, DNS, FTP, FTP_GET, FTP_PUT, HTTP, HTTPS, IMAP, MYSQL, NTP, PING, POP3, SMTP, SSH, SYSLOG, and TRACEROUTE.
4. Select **OK**.
5. Create another new **Service Group**.
6. Name the group OSPF_Services.
7. Add the following services to the *Members* list: DNS, FTP, FTP_GET, FTP_PUT, HTTP, HTTPS, IMAP, MYSQL, NTP, OSPF, PING, POP3, SMTP, SSH, SYSLOG, and TRACEROUTE.
8. Select **OK**.

To configure firewall service groups - CLI:

```
config firewall service group
  edit "BGP_services"
    set member "BGP", "DHCP" "DNS" "FTP" "FTP_GET" "FTP_PUT" "HTTP" "HTTPS" "IMAP"
      "MYSQL" "NTP" "PING" "POP3" "SMTP" "SSH" "TRACEROUTE" "SYSLOG"
  next
  edit "OSPF_services"
    set member "DHCP" "DNS" "FTP" "FTP_GET" "FTP_PUT" "HTTP" "HTTPS" "IMAP" "MYSQL"
      "NTP" "PING" "POP3" "SMTP" "SSH" "TRACEROUTE" "SYSLOG" "OSPF"
  next
end
```

Configuring BGP on the FortiGate

The only change from the standard BGP configuration for this example is configuring the blocking Peer 1's routes from being advertised to Peer 2. From the network topology you can guess that both of these peers likely share many routes in common and it doesn't make sense to advertise unneeded routes.

Blocking Peer 1's routes to Peer 2 is done with the `distribute-list-out` keyword. They allow you to select which routes you will advertise to a neighbor using an access list. In this case, we'll block all incoming routes from Peer 1 when we send updates to Peer 2. Otherwise Peer 1 and Peer 2 are regular neighbors.

The FortiGate redistributes routes learned from OSPF into BGP.

This is advanced configuration and the commands are only available in the CLI.

To create access list to block Peer 1 - CLI:

```
config access-list
  edit "block_peer1"
    config rule
      edit 1
        set prefix 172.21.111.0 255.255.255.0
        set action deny
        set exact-match enable
      end
    end
  end
```

To configure BGP on the FortiGate unit - CLI:

```
config router bgp
  set as 65001
  set router-id 10.11.201.110
  config redistribute ospf
    set status enable
  end
  config neighbor
    edit 172.22.222.5
      set remote-as 65001
      set distribute-list-out "block_peer1"
    next
    edit 172.21.111.5
      set remote-as 65001
    end
  end
end
```

Configuring OSPF on the FortiGate

This configuration involves only one OSPF area, so all traffic will be intra-area. If there were two or more areas with traffic going between them, it would be inter-area traffic. These two types are comparable to BGP's traffic within one AS (iBGP) or between multiple ASes (eBGP). Redistributing routes from OSPF to BGP is considered external because either the start or end point is a different routing protocol.

The OSPF configuration is basic, apart from redistributing BGP routes learned.

To configure OSPF on the FortiGate unit - GUI:

1. Go to **Router > Dynamic > OSPF**.
2. For Router ID enter `10.11.201.110` and then select **Apply**.
3. Under **Advanced Options > Redistribute**, select **BGP** and set the BGP **Metric** to 1.
4. For **Areas**, select **Create New**, enter the following information and then select **OK**.

Area (IP)	0.0.0.0
Type	Regular
Authentication	None

5. For **Networks**, select **Create New**.
6. Enter 10.11.201.0/255.255.255.0 for **IP/Netmask**, and select **OK**.
7. For **Interfaces**, select **Create New**.
8. Enter `OSPF_dmz_network` for **Name**.
9. Select `port1 (dmz)` for **Interface** and then select **OK**.

To configure OSPF on the FortiGate - CLI:

```
config router ospf
  set router-id 10.11.201.110
  config area
    edit 0.0.0.0
      set type regular
      set authentication none
    end
  config network
    edit 1
      set area 0.0.0.0
      set prefix 10.11.201.0 255.255.255.0
    end
  config interface
    edit "OSPF_dmz_network"
      set interface port1(dmz)
      set status enable
    end
  config redistribute bgp
    set status enable
    set metric 1
  end
end
```

Configuring other networking devices

As with all BGP configurations, the peer routers will need to be updated with the FortiGate device's BGP information, including IP address, AS number, and what capabilities are being used, such as IPv6, graceful restart, BFD, and so on.

Configuring ECMP support for BGP

Equal Cost Multiple Path (ECMP) is a mechanism that allows multiple routes to the same destination with different next-hops and load-balances routed traffic over those multiple next-hops.

- ECMP only works for routes that are sourced by the same routing protocol (Static Routes, OSPF, and BGP).
- ECMP is enabled, by default, with 10 paths.
- ECMP with static routes is effective if the routes are configured with the same distance and same priority.

To configure ECMP support - CLI:

```
config router bgp
    set ebgp-multipath disable[|enable]
    set ibgp-multipath disable[|enable]
    ...
end
```

Testing network configuration

Testing this configuration involves the standard connectivity checks, but also ensures that routes are being passed between protocols as expected.

Check the routing table on the FortiGate to ensure that routes from both OSPF and BGP are present.

Check the routing table on devices on the OSPF network for routes redistributed from BGP. Also, check those devices for connectivity to the Internet.

Check the routing table on Peer 2 to ensure that no routes from Peer 1 are present, but routes from the internal OSPF network are present.

For help with troubleshooting, see ["Troubleshooting BGP" on page 2201](#).

IS-IS

Intermediate System to Intermediate System Protocol (IS-IS) allows routing of ISO's OSI protocol stack Connectionless Network Service (CLNS). IS-IS is an Interior Gateway Protocol (IGP) that isn't intended to be used between Autonomous Systems (AS).

IS-IS was developed by Digital Equipment Corporation and later standardized by ISO in 1992 as ISO 19589 (see [RFC 1142](#), note that this RFC is different from the ISO version). About the same time, the Internet Engineering Task Force developed OSPF (see ["OSPF" on page 2146](#)). After the initial version, IP support was added to IS-IS and this version was called Integrated IS-IS (see [RFC 1195](#)). Its widespread use started when an early version of IS-IS was included with BSD v4.3 Linux as the routed daemon. The routing algorithm used by IS-IS, the Bellman–Ford algorithm, first saw widespread use as the initial routing algorithm of the ARPANET.

IS-IS is a link state protocol that is well-suited to smaller networks. It's in widespread use and has near universal support on routing hardware. It's quick to configure and works well if there are no redundant paths. However, IS-IS updates are sent out node-by-node, so it can be slow to find a path around network outages. IS-IS also lacks good authentication, can't choose routes based on different quality of service methods, and can create network loops if you're not careful. IS-IS uses Dijkstra's algorithm to find the best path, like OSPF.

While OSPF is more widely known, IS-IS is a viable alternative to OSPF in enterprise networks and ISP infrastructures, largely due to its native support for IPv6 and its non-disruptive methods for splitting, merging, migrating, and renumbering network areas.

FortiGate supports IS-IS for IPv4 and IPv6.

How IS-IS works

As one of the original modern dynamic routing protocols, IS-IS is straightforward. Its routing algorithm isn't complex, there are some options to allow fine-tuning, and it's straightforward to configure IS-IS on a FortiGate.

From [RFC 1142](#):

The routing algorithm used by the Decision Process is a shortest path first (SPF) algorithm.

Instances of the algorithm are run independently and concurrently by all intermediate systems in a routing domain. IntraDomain routing of a PDU occurs on a hop-by-hop basis: that is, the algorithm determines only the next hop, not the complete path, that a data PDU will take to reach its destination.

IS-IS versus static routing

IS-IS was one of the earliest dynamic routing protocols to work with IP addresses. As such, it's not as complex as more recent protocols. However, IS-IS is a big step forward from simple static routing.

While IS-IS may be slow in response to network outages, static routing has zero response. The same is true for convergence: static routing has zero convergence. Both IS-IS and static routing have a limited hop count, so it's not a strength or a weakness.

TLV

IS-IS uses type-length-value (TLV) parameters to carry information in Link-State PDUs (LSPs). Each IS-IS LSP consists of a variable-length header to which TLVs are appended in order to extend IS-IS for IP routing. The TLV field consists of one octet of type (T), one octet of length (L), and "L" octets of value (V). They're included in all of the IS-IS ["Packet types" on page 2227](#). For a complete breakdown of the LSP, see ["LSP structure" on page 2225](#).

In IS-IS, TLVs are used to determine route-leaking and authentication and are also used for IPv4 and IPv6 awareness and reachability.

- To determine which TLVs are responsible for route-leaking, see ["Default routing" on page 2227](#).
- To determine which TLVs are responsible for authentication, see ["Authentication" on page 2230](#).

For a complete list of reserved TLV codepoints, refer to [RFC 3359](#).

LSP structure

It's difficult to fully understand a routing protocol without knowing what information is carried in its packets. Knowing how routers exchange each type of information will help you better understand the IS-IS protocol and will allow you to configure your network more appropriately.

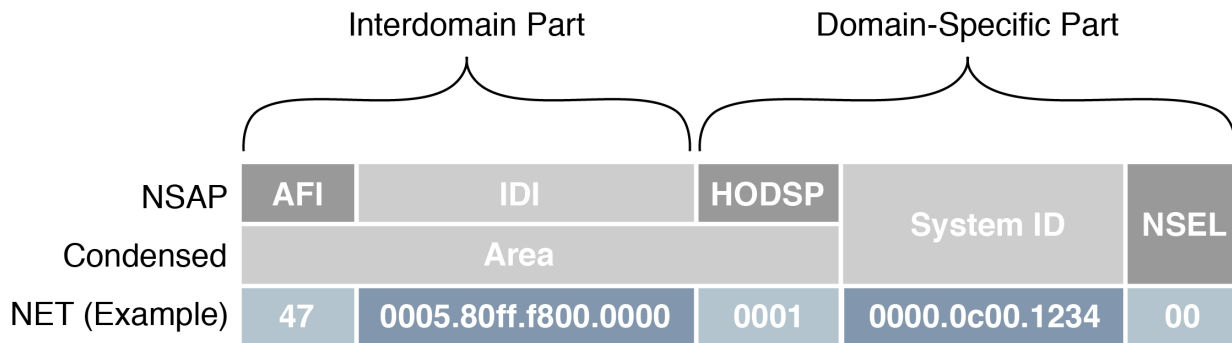
This section provides information about the contents of the IS-IS LSP. LSPs describe the network topology and can include IP routes and checksums.

NSAP and NET

IS-IS routing protocol utilizes ISO network addressing to identify network interfaces. The addresses are known as Network Service Access Points (NSAP). In general, IS-IS routers consist of only one NSAP, whereas IP addressing requires one IP address per interface.

In IS-IS, the NSAP address is translated into a Network Entity Title (NET), which is the same as the NSAP but can differentiate end systems by way of a byte called the n-selector (NSEL). In order for adjacencies to form in IS-IS, the NSEL must be set to zero, to indicate "this system". The total NET can be anywhere between 8 and 20 bytes long due to the support for variable length area addressing.

The following diagram identifies the individual parts of the NSAP, with explanations below:

NSAP and NET example

- **AFI** : The Authority and Format Identifier (AFI) specifies the format of the addressing family used. IS-IS is designed to carry routing information for several different protocols. Each entry has an address family identifier that identifies the globally unique Interdomain Part (IDP). For example, 49 is the AFI for private addresses, whereas 47 is the AFI for international organizations.
- **IDI**: The Initial Domain Identifier (IDI) identifies the routing domain within an interconnected network. The length of the IDI is typically determined by the AFI. If you are using an AFI of 49, you don't need to specify an IDI since the network is private.
- **HODSP** : The High Order Domain-Specific Part (HODSP) identifies the unique address within a specific routing domain. Together, the AFI, IDI, and HODSP define the area address. All of the nodes within an area must have the same area address.
- **System ID** : The *System ID* represents the 6-8 byte router identifier. The ID could be Media Access Control (MAC) format, as in the example above, or a static length IP address expressed in binary-coded decimal (BCD) format.
- **NSEL**: The *n-selector* (*NSEL*), as previously described, identifies the network layer transport service and must always be set to zero for IS-IS NETs.

Parts and terminology of IS-IS

Before you can understand how IS-IS functions, you need to understand some of the main concepts and parts of IS-IS.

DIS election and pseudonode LSP

In IS-IS routing protocol, a single router is chosen to be the designated intermediate system (DIS). The election of the DIS is determined automatically and dynamically on the LAN depending on highest interface priority and the subnetwork point of attachment (SNPA). The FortiGate is typically the DIS, and each router in its LAN is an intermediate system (IS).

Unlike OSPF, which elects a designated router (DR) and backup designated router (BDR), the DIS has no backup and determines the election of a new DIS whenever a router is added to the LAN or whenever the current DIS drops. A backup DIS is irrelevant since all of the routers on an IS-IS system are synchronized, and the short Hello interval used by the DIS quickly detects failures and the subsequent replacement of the DIS.

Synchronization of all the nodes in an IS-IS area could prove troublesome when updating the network infrastructure and would demand ever-increasing resources each time a new router is added (at an exponential scale). For this purpose, the DIS creates a pseudonode, which is essentially a virtual, logical node representing the LAN. The pseudonode requests adjacency status from all the routers in a multi-access network by sending IS-IS Hello (IIH) PDUs to Level 1 and Level 2 routers (where Level 1 routers share the same address as the DIS and

Level 2 routers do not). Using a pseudonode to alter the representation of the LAN in the link-state database (LSDB) greatly reduces the amount of adjacencies that area routers have to report. In essence, a pseudonode *collapses* a LAN topology, which allows a more linear scale to link-state advertising.

In order to maintain the database synchronization, the DIS periodically sends complete sequence number packets (CSNPs) to all participating routers.

Packet types

Four general packet types (PDUs) are communicated through IS-IS, appearing at both Level 1 and Level 2. They are described below.

- **Intermediate System-to-Intermediate System Hello (IIH) PDU** : As mentioned previously, the IIH PDU, or Hello packet, detects neighboring routers and indicates to the pseudonode the area's adjacency mesh. The Hello packet, flooded to the multicast address, contains the system ID of the sending router, the holding time, the circuit type of the interface on which the PDU was sent, the PDU length, the DIS identifier, and the interface priority (used in DIS election). The Hello packet also informs its area routers that it is the DIS. Hello packets are padded to the maximum IS-IS PDU size of 1492 bytes (the full MTU size) to assist in the detection of transmission errors with large frames or with MTU mismatches between adjacencies. The DIS typically floods Hello packets to the entire LAN every three seconds.
- **Link-state PDU (LSP)** : The LSP contains information about each router in an area and its connected interfaces. LSPs are refreshed periodically and acknowledged on the network by way of sequence number PDUs. If new LSP information is found, based on the most recent complete sequence number PDU (CSNP), out-of-date entries in the link-state database (LSDB) are removed and the LSDB is updated. For a more detailed breakdown of the LSP, see ["LSP structure" on page 2225](#).
- **Complete sequence number PDU (CSNP)**: CSNPs contain a list of all LSPs in the current LSDB. The CSNP informs other area routers of missing or outdated links in the adjacency mesh. The receiving routers then use this information to update their own database to ensure that all area routers converge. In contrast to Hello packets, CSNPs are sent every ten seconds and only between neighbors. In other words, they're never flooded.
- **Partial sequence number PDU (PSNP)** : PSNPs are used to request and acknowledge LSP information from an adjacency. When a router compares a CSNP with its local database and determines a discrepancy, the router requests an updated LSP using a PSNP. Once received, the router stores the LSP in its local database and responds to the DIS with acknowledgement.

Default routing

The default route is used if there are no other routes in the routing table or if none of the other routes apply to a destination. Including the gateway in the default route gives all traffic a next-hop address to use when leaving the local network. The gateway address is normally another router on the edge of the local network.

FortiGate units come with a default static route with an IPv4 address of 0.0.0.0, an administration distance of 10, and a gateway IPv4 address. Beginner administrators can use the default route settings until a more advanced configuration is warranted.

By default, all routes are displayed in the Routing Monitor list. To display the routes in the routing table, go to **Monitor > Routing Monitor**.

Route leaking

Route leaking is a term that's used to describe the bidirectional flow of information between internal and external routing interfaces. By default, IS-IS leaks routing information from a Level 1 area into a Level 2 area. In order to leak Level 2 routing information into a Level 1 area, you must configure an export policy. The ATT bit uses Type

Level Value (TLV) 128 (for internal reachability) and TLV 130 (for external IP address information) to determine whether or not a route is leaked. For more information about TLVs, see ["Troubleshooting IS-IS" on page 2231](#).

To configure IS-IS route leaking - CLI:

- On a Level 1-2 router:

```
config router isis
    set {redistribute-l2|redistribute6-l2} enable
end
```

- On a Level 1 router:

```
config router isis
    get router {info|info6} routing-table isis
    get router {info|info6} isis route
end
```

Default information originate

You can enable the default-information-originate option to generate and advertise a default route into the FortiGate device's IS-IS-enabled networks. The generated route may be based on routes that the FortiGate learns through a dynamic routing protocol, routes in the routing table, or both. IS-IS doesn't create the default route unless you use the `always` option.

If you experience any issues or if you wish to advertise your own static routes into IS-IS updates, set this to `disable`.

To enable the default information originate option - CLI:

```
config router isis
    set {default-originate | default-originate6}
end
```

Timer options

IS-IS uses various timers to regulate its performance, including garbage, update, and timeout timers. The FortiGate unit default timer settings (30, 180, and 120 seconds respectively) are effective in most configurations. If you change these settings, ensure that the new settings are compatible with local routers and access servers.

To configure the IS-IS timers - CLI:

```
config router isis
    set garbage-timer
    set update-timer
    set timeout-timer
end
```

You will find more information on each timer below.

Update timer

The update timer determines the interval between routing updates. Generally, this value is set to 30 seconds. There's some randomness added to help prevent network traffic congestion, which could result from all routers

simultaneously attempting to update their neighbors. The update timer should be at least three times smaller than the timeout timer, or you'll experience an error.

If you're experiencing significant traffic on your network, you can increase this interval to send fewer updates per minute. However, ensure you increase the interval for all the routers on your network or you'll experience timeouts that will degrade your network speed.

Timeout timer

The timeout timer is the maximum amount of time (in seconds) that a route is considered reachable while no updates are received for the route. This is the maximum time the DIS will keep a reachable route in the routing table while no updates for that route are received. If the DIS receives an update for the route before the timeout period expires, the timer is restarted. The timeout period should be at least three times longer than the update period, or you'll experience an error.

If you're experiencing problems with routers not responding in time to updates, increase this timer. However, remember that longer timeout intervals result in longer overall update periods. It may be a considerable amount of time before the DIS is done waiting for all the timers to expire on unresponsive routes.

Garbage timer

The garbage timer is the amount of time (in seconds) that the DIS will advertise a route as being unreachable before deleting the route from the routing table. If this timer is shorter, it will keep more up-to-date routes in the routing table and remove old ones faster. This results in a smaller routing table, which is useful if you have a very large network or if your network changes frequently.

IS-IS interface advertisements

You can use the `adv-passive-only` CLI command to configure which IS-IS interfaces the FortiGate advertises. If you set this command to enable, the FortiGate advertises only passive interfaces. If you set this to disable, the FortiGate advertises all IS-IS enabled interfaces.

To configure IS-IS interface advertisements - CLI:

```
config route isis
  set {adv-passive-only|adv-passive-only6} {enable | disable}
end
```

Loopback interfaces

You can configure loopback interfaces to run IS-IS.

To configure loopback interfaces to run IS-IS - CLI:

```
config router isis
  config isis-interface
    edit <name>
      set network-type loopback
    end
  end
end
```

Authentication

In routing protocols, it's typically desirable to establish authentication rules that prevent malicious and otherwise unwanted information from being injected into the routing table. IS-IS routing protocol utilizes TLV 10 to establish authentication. For more information about TLVs, see ["TLV" on page 2225](#).

Initially, IS-IS used plain cleartext to navigate the authentication rules, but this was found to be insecure since the cleartext packets were unencrypted and could be exposed to packet sniffers. As per [RFC 3567](#), HMAC-MD5 and enhanced cleartext authentication features were introduced in IS-IS, both of which encrypt authentication data, making them considerably more secure than using plain cleartext authentication.

HMAC-MD5 authentication

Hashed Message Authentication Codes - Message Digest 5 (HMAC-MD5) is a mechanism for applying a cryptographic hash function to the message authentication process. It is applied at both Level 1 and Level 2 routing. In IS-IS, an HMAC-MD5 can be applied to each type of LSP, on different interfaces, and with different passwords.

Authentication data is hashed using an AH (Authentication Header) key. From [RFC 2085](#):

The "AH Key" is used as a shared secret between two communicating parties. The Key is not a "cryptographic key" as used in a traditional sense. Instead, the AH key (shared secret) is hashed with the transmitted data and thus, assures that an intervening party cannot duplicate the authentication data. [...] Implementation should, and as frequently as possible, change the AH key. Keys need to be chosen at random, or generated using a cryptographically strong pseudo-random generator seeded with a random seed."

Cleartext authentication uses the configuration commands `area-password` and `domain-password` for authentication, but when migrating from cleartext authentication to HMAC-MD5, these command settings are automatically overwritten.

By the year 2005, the MD5 hash function had been identified as vulnerable to collision search attacks and various weaknesses. While such vulnerabilities don't compromise the use of MD5 within HMAC, administrators need to be aware of potential developments in cryptanalysis and cryptographic hash functions in the likely event that the underlying hash function needs to be replaced.

Enhanced cleartext authentication

Enhanced cleartext authentication is an extension to cleartext authentication that allows the encryption of passwords as they are displayed in the configuration. It includes a series of authentication mode commands and an authenticating key chain, and allows for more simple password modification and password management. Enhanced cleartext authentication also provides for smoother migration to and from changing authentication types. Intermediate systems continue to use the original authentication method until all the area routers are updated to use the new method.

Authentication key chain

A key chain is a list of one or more authentication keys including the send and receive lifetimes for each key. Keys are used for authenticating routing packets only during the specified lifetimes. A router migrates from one key to the next according to the scheduled send and receive lifetimes. If an active key is unavailable, the PDU is automatically discarded.

From [RFC 5310](#):

It should be noted that the cryptographic strength of the HMAC depends upon the cryptographic strength of the underlying hash function and on the size and quality of the key.

Troubleshooting IS-IS

Routing loops

Normally in routing, a path between two addresses is chosen and traffic is routed along that path from one address to the other. When there's a routing loop, that normal path doubles back on itself which creates a loop. When there are loops, the network has problems.

A routing loop occurs when a normally functioning network has an outage and one or more routers are offline. When packets encounter this, an alternate route is attempted to maneuver around the outage. During this phase it's possible for a route to be attempted that involves going back a hop and trying a different hop forward. If that hop forward is also blocked by the outage, a hop back and possibly the original hop forward may be selected. You can see if this continues, how it can consume not only network bandwidth but also many resources on the affected routers. The worst part is, this situation will continue until the network administrator changes the router settings or the downed routers come back online.

Routing loop effect on the network

In addition to this “traffic jam” of routed packets, every time the routing table for a router changes, that router sends an update out to all of the IS-IS routers connected to it. In a network loop, it's possible for a router to change its routes very quickly as it tries and fails along these new routes. This can quickly result in a flood of updates being sent out, which can effectively grind the network to a halt until the problem is fixed.

How to spot a routing loop

Any time network traffic slows down, you'll ask yourself if it's a network loop or not. Often slowdowns are normal. They're not a full stoppage and normal traffic resumes in a short period of time.

If the slowdown is a full halt of traffic, or a major slowdown doesn't return to normal quickly, you need to do serious troubleshooting quickly.

Some methods to troubleshoot your outage include:

- [Checking your logs](#)
- [Using SNMP network monitoring](#)
- [Using link health monitoring](#)
- [Looking at the packet flow](#)

If you're not running SNMP or link health monitoring, or if you have Fortinet routers that aren't Fortinet products in your network, you can use networking tools, such as ping and traceroute, to define the outage on your network and begin to fix it.

Checking your logs

If your routers log events to a central location, it can be easy to check the logs for your network for any outages.

On the FortiGate, go to **Log & Report > Log & Archive Access**. You'll want to look at both event logs and traffic logs. Events to look for will generally fall under CPU and memory usage, interfaces going offline (due to link health monitoring), and other similar system events.

Once you've found and fixed your network problem, you can go back to the logs and create a report to better see how things developed during the problem. This type of forensic analysis can better help you prepare for next time.

Using SNMP network monitoring

If your network had no problems one minute and slows to a halt the next, chances are something changed to cause that problem. Most of the time an offline router is the cause and once you find that router and bring it back online, things will return to normal.

If you can enable a hardware monitoring system such as SNMP or sFlow on your routers, you can be notified of the outage and where it's located, as soon as it happens.

Ideally you can configure SNMP on all your FortiGate routers and be alerted to all outages as they occur.

To use SNMP to detect potential routing loops - GUI:

1. Go to **System > Config > SNMP**.
2. Enable **SNMP Agent**.
3. Optionally, enter the **Description**, **Location**, and **Contact** information for this device for easier location of the problem report.
4. In either **SNMP v1/v2c** section or **SNMP v3** section, as appropriate, select **Create New**.
5. Enter the **Community Name** that you want to use.
6. In **Hosts**, select **Add** to add an IP address where you will be monitoring the FortiGate. You can add up to 8 different addresses.
7. Ensure that ports 161 and 162 (SNMP queries and traps) are allowed through your security policies.
8. In **SNMP Event**, select the events you want to be notified about. For routing loops, this should include **CPU Overusage**, **Memory Low**, and possibly **Log disk space low**. If there are problems, the log will fill up quickly, and the FortiGate device's resources will be overused.
9. Select **OK**.
10. Configure SNMP host (manager) software on your administration computer. This will monitor the SNMP information sent out by the FortiGate. Typically, you can configure this software to alert you about outages or CPU spikes that may indicate a routing loop.

Using link health monitoring

Another tool available to you on a FortiGate is the link health monitor. You can detect possible routing loops with link health monitors. You can configure the FortiGate to ping a gateway at regular intervals to ensure it's online and working. When the gateway isn't accessible, that interface is marked as down.

For more information about link health monitoring, see ["Link health monitor" on page 2025](#).

Looking at the packet flow

If you want to see what is happening on your network, look at the packets traveling on the network. In this situation, you're looking for routes that have metrics higher than 15, since that indicates that they're unreachable. Ideally, if you debug the flow of the packets and record the routes that are unreachable, you can create an accurate picture of the network outage.

Action to take on discovering a routing loop

Once you've mapped the problem on your network and determined it's in fact a routing loop, there are a number of steps you can take to correct it.

1. Get any offline routers back online. This may be a simple reboot or you may have to replace hardware. Often, this first step will restore your network to its normal operation, once the routing tables finish being updated.
2. Change your routing configuration on the edges of the outage. Even if step 1 brought your network back online, you should consider making changes to improve your network before the next outage occurs. These changes can include configuring features like holddowns and triggers for updates, split horizon, and poison reverse updates.

Split horizon and poison reverse updates

Split horizon is best explained with an example. You have three routers linked serially, let us call them routerA, routerB, and routerC. RouterA is linked only to routerB, routerC is linked only to routerB, and routerB is linked to both routerA and routerC. To get to routerC, routerA must go through routerB. If the link to routerC goes down, it's possible that routerB will try to use routerA's route to get to routerC. This route is A-B-C, so it won't work. However, if routerB tries to use it, this begins an endless loop. This situation is called a split horizon because from routerB's point of view, the horizon stretches out in each direction but in reality it only is on one side.

Poison reverse is the method used to prevent routes from running into split horizon problems. Poison reverse "poisons" routes away from the destination that use the current router in their route to the destination. This "poisoned" route is marked as unreachable for routers that can't use it. In IS-IS, this means that route is marked with a distance of 16.

Simple IS-IS example

This is an example of a typical medium-sized network configuration using IS-IS routing.

Imagine a company with four FortiGate devices connected to one another. A FortiGate at one end of the network connects to two routers, each with its own local subnet. One of these routers uses OSPF and the other router uses RIP.


Your task is to configure the four FortiGate devices to route traffic and process network updates using IS-IS, so that the farthest FortiGate (see 'FGT4' in ["Network layout and assumptions" on page 2234](#)) receives route updates for the two routers at the opposite end of the network. Furthermore, FGT4 has been given a loopback subnet that must be identified by the router running RIP.

Since the internal networks use OSPF and RIP, those protocols need to be redistributed throughout the IS-IS network. To keep the example simple, there will be no authentication of router traffic.


With IS-IS properly configured in this example, if a router fails or temporarily goes offline, the route change will propagate throughout the system.

Network layout and assumptions

Routing domains



IP scheme and interfaces



- It's assumed that each FortiGate is operating in NAT mode, running FortiOS 4.0MR2+.
- All interfaces have been previously assigned and no static routes are required.

- The Authority and Format Identifier (AFI) used is 49 : Locally administered (private).
- The Area identifiers are 0048 and 0049.

Expectations

- FGT4 must get the IS-IS route updates for RTR1 and RTR2 local subnets (10.1.1.0, 10.2.2.0, 10.3.3.0, 10.4.4.0).
- RTR1 must receive (via RIP2) the loopback subnet of FGT4 (10.60.60.1/32).

CLI configuration

The following CLI configuration occurs on each FortiGate (as identified), including only the relevant parts.

FGT1

```
config router isis
  config isis-interface
    edit "port3"
      set circuit-type level-1
      set network-type broadcast
      set status enable
    next
  end
  config isis-net
    edit 1
      set net 49.0048.1921.6818.2136.00
    next
  end
  config redistribute "connected"
  end
  config redistribute "rip"
    set status enable
    set level level-1
  end
  config redistribute "ospf"
    set status enable
    set level level-1
  end
end
config router rip
  config interface
    edit "port2"
      set receive-version 2
      set send-version 2
    next
  end
  config network
    edit 1
      set prefix 10.10.10.0 255.255.255.0
    next
  end
  config redistribute "isis"
    set status enable
  end
end
```

FGT2

```
config router isis
config isis-interface
edit "port3"
set circuit-type level-1
set network-type broadcast
set status enable
next
edit "port2"
set network-type broadcast
set status enable
next
end
config isis-net
edit 1
set net 49.0048.1221.6818.2110.00
next
end
set redistribute-l1 enable
set redistribute-l2 enable
end
```

FGT3

```
config router isis
set is-type level-2-only
config isis-interface
edit "wan1"
set network-type broadcast
set status enable
next
edit "dmz1"
set network-type broadcast
set status enable
next
end
config isis-net
edit 1
set net 49.0048.1921.6818.2108.00
next
edit 2
set net 49.0049.1921.6818.2108.00
next
end
end
```

FGT4

```
config router isis
set is-type level-2-only
config isis-interface
edit "wan1"
set network-type broadcast
set status enable
```

```

        next
    end
    config isis-net
        edit 1
            set net 49.0049.1721.0160.1004.00
        next
    end
    config redistribute "connected"
        set status enable
    end
end

```

Verification

Once the network has been configured, you need to test that it works as expected. Use the following CLI commands on the devices indicated.

Verifying if RTR1 receives loopback subnet of FGT4

```
(RTR1) # get router info routing-table all
```

Result:

```

C    10.1.1.0/24 is directly connected, vlan1
C    10.2.2.0/24 is directly connected, vlan2
C    10.10.10.0/24 is directly connected, dmz1
R    10.40.40.0/24 [120/2] via 10.10.10.1, dmz1, 00:04:07
R    10.50.50.0/24 [120/2] via 10.10.10.1, dmz1, 00:04:07
R    10.60.60.1/32 [120/2] via 10.10.10.1, dmz1, 00:04:07

```

(*) If required, filtering out 10.50.50.0 and 10.40.40.0 from the routing table could be done with a route-map.

Verification on FGT2, which is the border between L1 and L2 routing levels, looking at IS-IS information

```
FGT2 # get router info isis interface
```

Result:

```

port2 is up, line protocol is up
Routing Protocol: IS-IS ((null))
Network Type: Broadcast
Circuit Type: level-1-2
Local circuit ID: 0x01
Extended Local circuit ID: 0x00000003
Local SNPA: 0009.0f85.ad8c
IP interface address:
10.40.40.2/24
IPv4 interface address:
Level-1 Metric: 10/10, Priority: 64, Circuit ID: 1221.6818.2110.01
Number of active level-1 adjacencies: 0
Level-2 Metric: 10/10, Priority: 64, Circuit ID: 1221.6818.2110.01
Number of active level-2 adjacencies: 1
Next IS-IS LAN Level-1 Hello in 6 seconds
Next IS-IS LAN Level-2 Hello in 1 seconds
port3 is up, line protocol is up
Routing Protocol: IS-IS ((null))

```

```

Network Type: Broadcast
Circuit Type: level-1
Local circuit ID: 0x02
Extended Local circuit ID: 0x00000004
Local SNPA: 0009.0f85.ad8d
IP interface address:
10.30.30.2/24
IPv4 interface address:
Level-1 Metric: 10/10, Priority: 64, Circuit ID: 1221.6818.2110.02
Number of active level-1 adjacencies: 1
Next IS-IS LAN Level-1 Hello in 2 seconds

```

```
FGT2 # get router info isis neighbor
```

Result:

System Id	Interface	SNPA	State	Holdtime	Type	Protocol
1921.6818.2108	port2	0009.0f04.0794	Up	22	L2	IS-IS
1921.6818.2136	port3	0009.0f85.acf7	Up	29	L1	IS-IS

Verification on FGT3, which is border between 2 areas, looking at IS-IS information

IS-IS router CLI commands available:

```
FGT3 # get router info isis ?
```

Result:

interface	show isis interfaces
neighbour	show CLNS neighbor adjacencies
is-neighbour	show IS neighbor adjacencies
database	show IS-IS link state database
route	show IS-IS IP routing table
topology	show IS-IS paths

Example of interface status and neighbors:

```
FGT3 # get router info isis interface
```

Result:

```

wan1 is up, line protocol is up
Routing Protocol: IS-IS ((null))
Network Type: Broadcast
Circuit Type: level-1-2
Local circuit ID: 0x01
Extended Local circuit ID: 0x00000003
Local SNPA: 0009.0f04.0794

```

```

IP interface address:
10.40.40.1/24
IPv4 interface address:
Level-2 Metric: 10/10, Priority: 64, Circuit ID: 1221.6818.2110.01
Number of active level-2 adjacencies: 1
Next IS-IS LAN Level-2 Hello in 3 seconds

```

```

dmz1 is up, line protocol is up
Routing Protocol: IS-IS ((null))
Network Type: Broadcast
Circuit Type: level-1-2
Local circuit ID: 0x02
Extended Local circuit ID: 0x00000005
Local SNPA: 0009.0f04.0792
IP interface address:
10.50.50.1/24
IPv4 interface address:
Level-2 Metric: 10/10, Priority: 64, Circuit ID: 1721.0160.1004.01
Number of active level-2 adjacencies: 1
Next IS-IS LAN Level-2 Hello in 7 seconds

```

```
FGT3 # get router info isis neighbor
```

Result:

System Id	Interface	SNPA	State	Holdtime	Type	Protocol
1221.6818.2110	wan1	0009.0f85.ad8c	Up	8	L2	IS-IS
1721.0160.1004	dmz1	0009.0f52.7704	Up	8	L2	IS-IS

Verification on FGT4 that the remote subnets from RTR1 and RTR2 are in the routing table and learned with IS-IS

```
FGT4 # get router info routing-table all
```

Result:

```

Codes: K - kernel, C - connected, S - static, R - RIP, B - BGP
O - OSPF, IA - OSPF inter area
  N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
  E1 - OSPF external type 1, E2 - OSPF external type 2
  i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
  * - candidate default
i L2 10.1.1.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
  i L2 10.2.2.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
  i L2 10.3.3.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
  i L2 10.4.4.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
i L2 10.10.10.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
  i L2 10.11.11.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
  i L2 10.20.20.0/24 [115/30] via 10.50.50.1, wan1, 00:12:46
  i L2 10.30.30.0/24 [115/30] via 10.50.50.1, wan1, 00:13:55
  i L2 10.40.40.0/24 [115/20] via 10.50.50.1, wan1, 00:15:30
C 10.50.50.0/24 is directly connected, wan1
C 10.60.60.1/32 is directly connected, loopback

```

Troubleshooting

The following diagnose commands are available for further IS-IS troubleshooting and will display all IS-IS activity (sent and received packets):

```
FGT # diagnose ip router isis level info
FGT # diagnose ip router isis all enable
FGT # diagnose debug enable
```

...to stop the debug type output:

```
FGT # diagnose ip router isis level none
```

Output and interpretation depends on the issue faced. You can provide this information to TAC if you open a support ticket.

Multicast forwarding

Multicasting (also called IP multicasting) consists of using a single multicast source to send data to many receivers. Multicasting can be used to send data to many receivers simultaneously while conserving bandwidth and reducing network traffic. Multicasting can be used for one-way delivery of media streams to multiple receivers and for one-way data transmission for news feeds, financial information, and so on.

Also, RIPv2 uses multicasting to share routing table information, OSPF uses multicasting to send hello packets and routing updates, Enhanced Interior Gateway Routing Protocol (EIGRP) uses multicasting to send routing information to all EIGRP routers on a network segment and the Bonjour network service uses multicasting for DNS.

A FortiGate can operate as a Protocol Independent Multicast (PIM) version 2 router. FortiGate devices support PIM sparse mode (RFC 4601) and PIM dense mode (RFC 3973) and can service multicast servers or receivers on the network segment to which a FortiGate interface is connected. Multicast routing isn't supported in transparent mode (TP mode).



To support PIM communications, the sending and receiving applications and all connecting PIM routers in between must be enabled with PIM version 2. PIM can use static routes, RIP, OSPF, or BGP to forward multicast packets to their destinations. To enable source-to-destination packet delivery, either sparse mode or dense mode must be enabled on the PIM-router interfaces. Sparse mode routers can't send multicast messages to dense mode routers. In addition, if a FortiGate is located between a source and a PIM router, two PIM routers, or is connected directly to a receiver, you must create a security policy manually to pass encapsulated (multicast) packets or decapsulated data (IP traffic) between the source and destination.

A PIM domain is a logical area comprising a number of contiguous networks. The domain contains at least one Boot Strap Router (BSR), and if sparse mode is enabled, a number of Rendezvous Points (RPs) and Designated Routers (DRs). When PIM is enabled on a FortiGate, the FortiGate can perform any of these functions at any time as configured.

Sparse mode

Initially, all candidate BSRs in a PIM domain exchange bootstrap messages to select one BSR to which each RP sends the multicast address or addresses of the multicast groups that it can service. The selected BSR chooses one RP per multicast group and makes this information available to all of the PIM routers in the domain through bootstrap messages. PIM routers use the information to build packet distribution trees, which map each multicast group to a specific RP. Packet distribution trees may also contain information about the sources and receivers associated with particular multicast groups.



When a FortiGate interface is configured as a multicast interface, sparse mode is enabled on it by default to ensure that distribution trees are not built unless at least one downstream receiver requests multicast traffic from a specific source. If the sources of multicast traffic and their receivers are close to each other and the PIM domain contains a dense population of active receivers, you may choose to enable dense mode throughout the PIM domain instead.

An RP represents the root of a non-source-specific distribution tree to a multicast group. By joining and pruning the information contained in distribution trees, a single stream of multicast packets (for example, a video feed) originating from the source can be forwarded to a certain RP to reach a multicast destination.

Each PIM router maintains a Multicast Routing Information Base (MRIB) that determines to which neighboring PIM router join and prune messages are sent. An MRIB contains reverse-path information that reveals the path of a multicast packet from its source to the PIM router that maintains the MRIB.

To send multicast traffic, a server application sends IP traffic to a multicast group address. The locally elected DR registers the sender with the RP that is associated with the target multicast group. The RP uses its MRIB to forward a single stream of IP packets from the source to the members of the multicast group. The IP packets are replicated only when necessary to distribute the data to branches of the RP distribution tree.

To receive multicast traffic, a client application can use Internet Group Management Protocol (IGMP) version 1 (RFC 1112), 2 (RFC 2236), or 3 (RFC 3376) control messages to request the traffic for a particular multicast group. The locally elected DR receives the request and adds the host to the multicast group that's associated with the connected network segment by sending a join message towards the RP for the group. Afterward, the DR queries the hosts on the connected network segment continually to determine whether the hosts are active. When the DR no longer receives confirmation that at least one member of the multicast group is still active, the DR sends a prune message towards the RP for the group.

FortiOS supports PIM sparse mode multicast routing for IPv6 multicast (multicast6) traffic and is compliant with RFC 4601: Protocol Independent Multicast - Sparse Mode (PIM-SM). You can use the following CLI commands to configure IPv6 PIM sparse multicast routing:

```
config router multicast6
  set multicast-routing {enable | disable}
  config interface
    edit <interface-name>
      set hello-interval <1-65535 seconds>
      set hello-holdtime <1-65535 seconds>
    end
  config pim-sm-global
    config rp-address
      edit <index>
        set ipv6-address <ipv6-address>
      end
```

The following diagnose commands for IPv6 PIM sparse mode are also available:

```
diagnose ipv6 multicast status
diagnose ipv6 multicast vif
diagnose ipv6 multicast mroute
```

Dense mode

The packet organization used in sparse mode is also used in dense mode. When a multicast source begins to send IP traffic and dense mode is enabled, the closest PIM router registers the IP traffic from the multicast source (S) and forwards multicast packets to the multicast group address (G). All PIM routers initially broadcast the multicast packets throughout the PIM domain to ensure that all receivers that have requested traffic for multicast group address G can access the information, if needed.

To forward multicast packets to specific destinations afterward, the PIM routers build distribution trees based on the information in multicast packets. Upstream PIM routers depend on prune/graft messages from downstream PIM routers to determine if receivers are actually present on directly-connected network segments. The PIM

routers exchange state refresh messages to update their distribution trees. FortiGate devices store this state information in a Tree Information Base (TIB), which is used to build a multicast forwarding table. The information in the multicast forwarding table determines whether packets are forwarded downstream. The forwarding table is updated whenever the TIB is modified.

PIM routers receive data streams every few minutes and update their forwarding tables using the source (S) and multicast group (G) information in the data stream. Superfluous multicast traffic is stopped by PIM routers that don't have downstream receivers. PIM routers that don't manage multicast groups send prune messages to the upstream PIM routers. When a receiver requests traffic for multicast address G, the closest PIM router sends a graft message upstream to begin receiving multicast packets.

FortiGate devices operating in NAT mode can also be configured as multicast routers. You can configure a FortiGate to be a Protocol Independent Multicast (PIM) router operating in Sparse Mode (SM) or Dense Mode (DM).

PIM support

You can configure a FortiGate to support PIM by going to **Network > Multicast** in the FortiGate GUI and enabling multicast routing. You can also enable multicast routing using the `config router multicast` CLI command. When PIM is enabled, the FortiGate allocates memory to manage mapping information. The FortiGate communicates with neighboring PIM routers to acquire mapping information and, if required, processes the multicast traffic associated with specific multicast groups.



The end-user multicast client-server applications must be installed and configured to initiate Internet connections and handle broadband content such as audio and video information.

Client applications send multicast data by registering IP traffic with a PIM-enabled router. An end user can type in a class D multicast group address, an alias for the multicast group address, or a conference call number to initiate the session.

Rather than sending multiple copies of generated IP traffic to more than one specific IP destination address, PIM-enabled routers encapsulate the data and use the one multicast group address to forward multicast packets to multiple destinations. Because one destination address is used, a single stream of data can be sent. Client applications receive multicast data by requesting that the traffic destined for a certain multicast group address be delivered to them. End users can use phone books, a menu of ongoing or future sessions, or some other method through a user interface to select the address of interest.

A class D address in the 224.0.0.0 to 239.255.255.255 range may be used as a multicast group address, subject to the rules assigned by the Internet Assigned Numbers Authority (IANA). All class D addresses must be assigned in advance. Because there isn't a way to determine in advance if a certain multicast group address is in use, collisions may occur (to resolve this problem, end users may switch to a different multicast address).

To configure a PIM domain

1. If you'll be using sparse mode, determine appropriate paths for multicast packets.
2. Make a note of the interfaces that will be PIM-enabled. These interfaces may run a unicast routing protocol.
3. If you'll be using sparse mode and want multicast packets to be handled by specific (static) RPs, record the IP addresses of the PIM-enabled interfaces on those RPs.
4. Enable PIM version 2 on all participating routers between the source and receivers. On FortiGate devices, use the `config router multicast` command to set global operating parameters.

5. Configure the PIM routers that have good connections throughout the PIM domain to be candidate BSRs.
6. If sparse mode is enabled, configure one or more of the PIM routers to be candidate RPs.
7. If required, adjust the default settings of PIM-enabled interfaces.

Multicast forwarding and FortiGate devices

In both transparent mode and NAT mode, you can configure FortiGate devices to forward multicast traffic.

For a FortiGate to forward multicast traffic, you must add FortiGate multicast security policies. Basic multicast security policies accept any multicast packets at one FortiGate interface and forward the packets out another FortiGate interface. You can also use multicast security policies to be selective about the multicast traffic that's accepted, based on source and destination address, and to perform NAT on multicast packets.

In the example shown below, a multicast source on the marketing network with IP address 192.168.5.18 sends multicast packets to the members of network 239.168.4.0. At the FortiGate, the source IP address for multicast packets originating from workstation 192.168.5.18 is translated to 192.168.18.10. In this example, the FortiGate isn't acting as a multicast router.

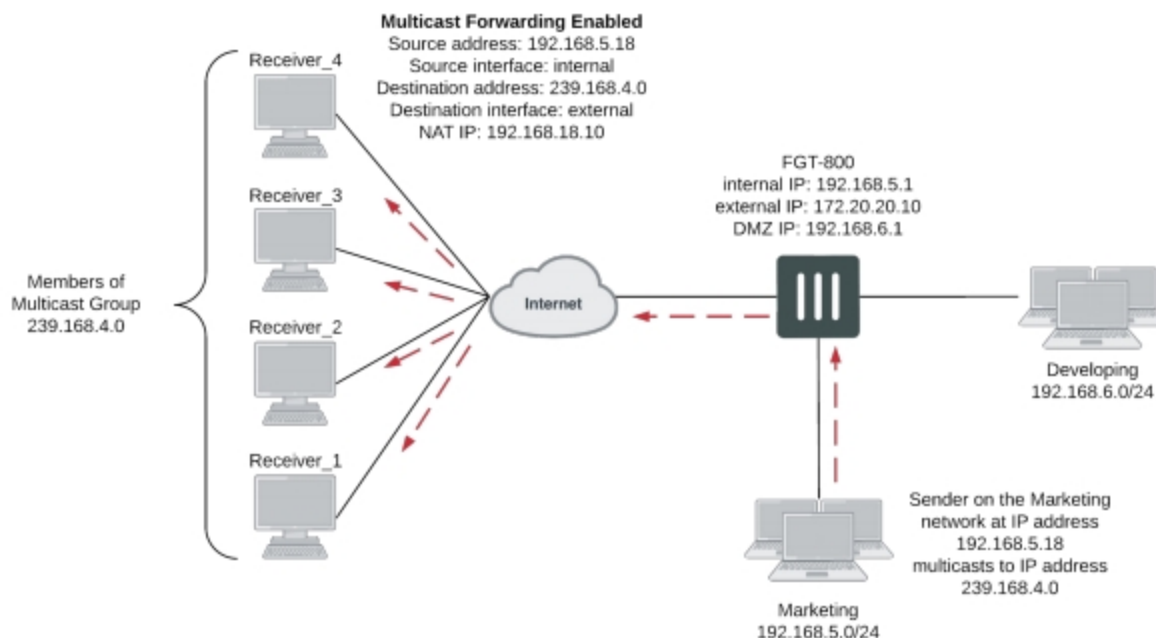
Multicast forwarding and RIPv2

RIPv2 uses multicast to share routing table information. If a FortiGate is installed on a network that includes RIPv2 routers, you must configure the FortiGate to forward multicast packets so that RIPv2 devices can share routing data through the FortiGate. No special FortiGate configuration is required to share RIPv2 data, you can simply use the information in the following sections to configure the FortiGate to forward multicast packets.



RIPv1 uses broadcasting to share routing table information. To allow RIPv1 packets through a FortiGate, you can add standard security policies. Security policies to accept RIPv1 packets can use the ANY predefined firewall service or the RIP predefined firewall service.

Example multicast network including a FortiGate that forwards multicast packets



Configuring FortiGate multicast forwarding

You configure FortiGate multicast forwarding in the CLI. Two steps are required:

1. Add multicast security policies
2. Enable multicast forwarding

This second step is required only if a FortiGate is operating in NAT mode. If the FortiGate is operating in transparent mode, adding a multicast policy enables multicast forwarding.



There's sometimes confusion between the terms “forwarding” and “routing”. These two functions shouldn't be taking place at the same time.

It's mentioned that multicast-forward should be enabled when the FortiGate is in NAT mode and that this will forward any multicast packet to all interfaces. However, this parameter shouldn't be enabled when the FortiGate operates as a multicast router (for example, with a routing protocol enabled). It should only be enabled when there's no routing protocols activated.

Adding multicast security policies

You need to add security policies to allow packets to pass from one interface to another. Multicast packets require multicast security policies. You add multicast security policies from the CLI using the `config firewall multicast-policy` command. As with unicast security policies, you specify the source and destination interfaces and, optionally, the allowed address ranges for the source and destination addresses of the packets.

You can also use multicast security policies to configure source NAT and destination NAT for multicast packets.

Keep the following in mind when configuring multicast security policies:

- The matched forwarded (outgoing) IP multicast source IP address is changed to the configured IP address.
- Source and destination interfaces are optional. If left blank, the multicast will be forwarded to ALL interfaces.
- Source and destination addresses are optional. If left unset, it means ALL addresses.
- The `nat` keyword is optional. Use it when source address translation is needed.

Enabling multicast forwarding

Multicast forwarding is enabled by default. In NAT mode you must use the `multicast-forward` keyword of the `system settings` CLI command to enable or disable multicast forwarding. When `multicast-forward` is enabled, a FortiGate forwards any multicast IP packets in which the TTL is 2 or higher to all interfaces and VLAN interfaces except the receiving interface. The TTL in the IP header will be reduced by 1. Even though the multicast packets are forwarded to all interfaces, you must add security policies to actually allow multicast packets through the FortiGate. In our example, the security policy allows multicast packets received by the internal interface to exit to the external interface.



Enabling multicast forwarding is only required if a FortiGate is operating in NAT mode. If a FortiGate unit is operating in transparent mode, adding a multicast policy enables multicast forwarding.

To enable multicast forwarding - CLI:

```
config system settings
set multicast-forward enable
```

```
end
```

If multicast forwarding is disabled and the FortiGate drops packets that have multicast source or destination addresses.

You can also use the `multicast-ttl-notchange` keyword of the `system settings` command so that the FortiGate doesn't increase the TTL value for forwarded multicast packets. You should use this option only if packets are expiring before reaching the multicast router.

```
config system settings
    set multicast-ttl-notchange enable
end
```

In transparent mode, a FortiGate doesn't forward frames with multicast destination addresses. Multicast traffic, such as the one used by routing protocols or streaming media, may need to traverse the FortiGate and shouldn't interfere with the communication. To avoid any issues during transmission, you can set up multicast security policies. These types of security policies can only be enabled using the CLI.



When you use multicast security policies, you must disable the `multicast-skip-policy` CLI parameter. To disable enter the commands:

```
config system settings
    set multicast-skip-policy disable
end
```

In this simple example, a check isn't performed on the source or destination interfaces. A multicast packet received on an interface is flooded unconditionally to all interfaces on the forwarding domain, except the incoming interface.

To enable the multicast policy - CLI:

```
config firewall multicast-policy
    edit 1
        set action accept
    end
```

In this example, the multicast policy only applies to the source port of WAN1 and the destination port of Internal.

To enable the restrictive multicast policy - CLI:

```
config firewall multicast-policy
    edit 1
        set srcintf wan1
        set dstintf internal
        set action accept
    end
```

In this example, packets are allowed to flow from WAN1 to Internal, and sourced by the address 172.20.120.129, which is represented by the address object "example_addr-1".

To enable the restrictive multicast policy - CLI:

```
config firewall multicast-policy
    edit 1
        set srcintf wan1
        set srcaddr example_addr-1
        set dstintf internal
        set action accept
```

```
end
```

This example shows how to configure the multicast security policy required for the configuration shown. This policy accepts multicast packets that are sent from a PC with IP address 192.168.5.18 to destination address range 239.168.4.0. The policy allows the multicast packets to enter the internal interface and then exit the external interface. When the packets leave the external interface, their source address is translated to 192.168.18.10

```
config firewall multicast-policy
edit 5
    set srcaddr 192.168.5.18 255.255.255.255
    set srcintf internal
    set destaddr 239.168.4.0 255.255.255.0
    set dstintf external
    set nat 192.168.18.10
end
```

This example shows how to configure a multicast security policy so that the FortiGate forwards multicast packets from a multicast server with an IP 10.10.10.10 is broadcasting to address 225.1.1.1. This server is on the network connected to the FortiGate DMZ interface.

```
config firewall multicast-policy
edit 1
    set srcintf DMZ
    set srcaddr 10.10.10.10 255.255.255.255
    set dstintf Internal
    set dstaddr 225.1.1.1 255.255.255.255
    set action accept
edit 2
    set action deny
end
```

Displaying IPv6 multicast router information

You can use the following CLI command to display IPv6 multicast router information (equivalent to the IPv4 version of the command):

```
get router info6 multicast
```


Multicast routing examples

This section contains the following multicast routing configuration examples and information:

- Example FortiGate PIM-SM configuration using a static RP
- FortiGate PIM-SM debugging examples
- Example multicast destination NAT (DNAT) configuration
- Example PIM configuration that uses BSR to find the RP

Example FortiGate PIM-SM configuration using a static RP

The example Protocol Independent Multicast Sparse Mode (PIM-SM) configuration shown below has been tested for multicast interoperability using PIM-SM between Cisco 3750 switches running 12.2 and a FortiGate-800 running FortiOS v3.0 MR5 patch 1. In this configuration, the receiver receives the multicast stream when it joins the group 233.254.200.1.

Example: FortiGate PIM-SM topology

The configuration uses a statically configured rendezvous point (RP) which resides on the Cisco_3750_1. Using a bootstrap router (BSR) wasn't tested in this example. See "Example PIM configuration that uses BSR to find the RP" for an example that uses a BSR.

Configuration steps

The following procedures show how to configure the multicast configuration settings for the devices in the example configuration.

- Cisco_3750_1 router configuration
- Cisco_3750_2 router configuration
- To configure the FortiGate-800 unit
- Cisco_3750_3 router configuration

Cisco_3750_1 router configuration

```
version 12.2
!
hostname Cisco-3750-1
!
switch 1 provision ws-c3750-24ts
ip subnet-zero
ip routing
!
ip multicast-routing distributed
!
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
interface Loopback0
    ip address 169.254.100.1 255.255.255.255
!
interface FastEthernet1/0/23
    switchport access vlan 182
    switchport mode access
!
interface FastEthernet1/0/24
    switchport access vlan 172
    switchport mode access
!
interface Vlan172
    ip address 10.31.138.1 255.255.255.0
    ip pim sparse-mode
    ip igmp query-interval 125
    ip mroute-cache distributed
!
interface Vlan182
    ip address 169.254.82.250 255.255.255.0
    ip pim sparse-mode
    ip mroute-cache distributed
!
ip classless
ip route 0.0.0.0 0.0.0.0 169.254.82.1
ip http server
ip pim rp-address 169.254.100.1 Source-RP
!
ip access-list standard Source-RP
    permit 233.254.200.0 0.0.0.255
```

Cisco_3750_2 router configuration

```
version 12.2
!
hostname Cisco-3750-2
!
switch 1 provision ws-c3750-24ts
ip subnet-zero
ip routing
!
ip multicast-routing distributed
```

```

!
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
interface FastEthernet1/0/23
    switchport access vlan 138
    switchport mode access
!
interface FastEthernet1/0/24
    switchport access vlan 182
    switchport mode access
!
interface Vlan138
    ip address 10.31.138.250 255.255.255.0
    ip pim sparse-mode
    ip mroute-cache distributed
!
interface Vlan182
    ip address 169.254.82.1 255.255.255.0
    ip pim sparse-mode
    ip mroute-cache distributed
!
ip classless
ip route 0.0.0.0 0.0.0.0 10.31.138.253
ip route 169.254.100.1 255.255.255.255 169.254.82.250
ip http server
ip pim rp-address 169.254.100.1 Source-RP
!
!
ip access-list standard Source-RP
permit 233.254.200.0 0.0.0.255

```

To configure the FortiGate-800 unit - GUI:

1. Configure the internal and external interfaces.

- **Internal**

Go to **Network > Interfaces**.

Select the internal interface.

Verify the following settings:

Type:	Physical Interface
Addressing mode:	Manual
IP/Network Mask:	10.31.138.253 255.255.255.0
Administrative Access:	PING

Select **OK**.

- **External**

Go to **Network > Interfaces**.
Select the external interface.

Verify the following settings:

Type:	Physical Interface
Addressing mode:	Manual
IP/Network Mask:	10.31.130.253 255.255.255.0
Administrative Access:	HTTPS and PING

Select **OK**.

2. Add a firewall addresses.

Go to **Policy & Objects > Addresses**.

- RP

Select **Create New**.

Use the following settings:

Category	Address
Name	RP
Type	Subnet
Subnet/IP Range	169.254.100.1/32
Interface	Any
Visibility	<enabled>

Select **OK**.

- Multicast source subnet

Select **Create New**.

Use the following settings:

Category	Address
Name	multicast_source_subnet

Type	Subnet
Subnet/IP Range	169.254.82.0/24
Interface	Any
Visibility	<enabled>

Select **OK**.

3. Add destination multicast address

Go to **Policy & Objects > Addresses**.

Select **Create New**.

Use the following settings:

Category	Multicast Address
Name	Multicast_stream
Type	Broadcast Subnet
Broadcast Subnet	233.254.200.0/24
Interface	Any
Visibility	<enabled>

Select **OK**.

4. Add standard security policies to allow traffic to reach the RP.

Go to **Policy & Objects > IPv4 Policy**.

- 1st policy

Select **Create New**.

Use the following settings:

Incoming Interface	internal
Source Address	all
Outgoing Interface	external
Destination Address	RP
Schedule	always

Service	ALL
Action	ACCEPT

Select **OK**.

- 2nd policy

Select **Create New**

Use the following settings:

Incoming Interface	external
Source Address	RP
Outgoing Interface	internal
Destination Address	all
Schedule	always
Service	ALL
Action	ACCEPT

Select **OK**.

5. Add the multicast security policy.

Go to **Policy & Objects > Multicast Policy**.

Select **Create New**.

Use the following settings:

Incoming Interface	external
Source Address	multicast_source_subnet
Outgoing Interface	internal
Destination Address	multicast_stream
Protocol	Any
Action	ACCEPT

Select **OK**.

6. Add an access list. (CLI only)

```
config router access-list
```

```

edit Source-RP
  config rule
    edit 1
      set prefix 233.254.200.0 255.255.255.0
      set exact-match disable
    next
  end

```

7. Add some static routes.

Go to **Network > Static Routes**.

- Route 1

Select **Create New**.

Use the following settings:

Destination IP/Mask	0.0.0.0/0.0.0.0
Gateway	10.31.130.250
Interface	internal
Administrative Distance	<default>
Priority	<default>

Select **OK**.

- Route 2

Select **Create New**.

Use the following settings:

Destination IP/Mask	169.254.0.0/16
Gateway	10.31.138.250
Interface	external
Administrative Distance	<default>
Priority	<default>

Select **OK**.

8. Configure multicast routing.

Go to **Network > Multicast**.

Add the following Static Rendezvous Point(s):

- 169.254.100.1
- Route 1

Select **Create New**.

Use the following settings:

Interface	internal
PIM Mode	Sparse Mode
DR Priority	<not needed in this scenario>
RP Candidate	<not needed in this scenario>
RP Candidate Priority	<not needed in this scenario>

Select **OK**.

- Route 2

Select **Create New**.

Use the following settings:

Interface	external
PIM Mode	Sparse Mode
DR Priority	
RP Candidate	
RP Candidate Priority	

Select **OK**.

Cisco_3750_3 router configuration

```

version 12.2
!
hostname Cisco-3750-3
!
switch 1 provision ws-c3750-24ts
ip subnet-zero
ip routing
!
ip multicast-routing distributed
!
spanning-tree mode pvst
no spanning-tree optimize bpdu transmission
spanning-tree extend system-id
!
interface FastEthernet1/0/23
    switchport access vlan 128
    switchport mode access
!
interface FastEthernet1/0/24

```

```

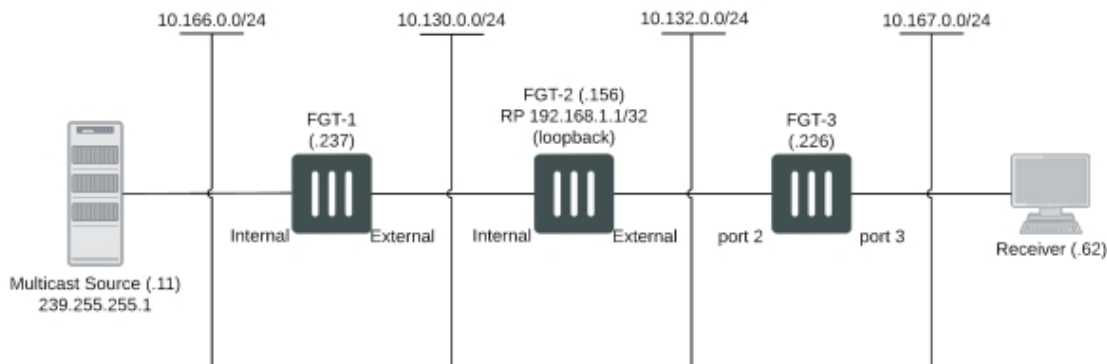
        switchport access vlan 130
        switchport mode access
    !
interface Vlan128
    ip address 10.31.128.130 255.255.255.252
    ip pim sparse-mode
    ip mroute-cache distributed
    !
interface Vlan130
    ip address 10.31.130.250 255.255.255.0
    ip pim sparse-mode
    ip mroute-cache distributed
    !
ip classless
ip route 0.0.0.0 0.0.0.0 10.31.130.1
ip http server
ip pim rp-address 169.254.100.1 Source-RP
    !
    !
ip access-list standard Source-RP
permit 233.254.200.0 0.0.0.255

```

FortiGate PIM-SM debugging examples

Using the example topology shown below, you can trace the multicast streams and states within the three FortiGate devices (FGT-1, FGT-2, and FGT-3) using the debug commands described in this section. The command output in this section is taken from the FortiGate when the multicast stream is flowing correctly from source to receiver.

PIM-SM debugging topology



Checking that the receiver has joined the required group

From the last hop router, FGT-3, you can use the following command to check that the receiver has correctly joined the required group.

```

FGT-3 # get router info multicast igmp groups
IGMP Connected Group Membership
Group Address Interface Uptime Expires Last Reporter
239.255.255.1 port3 00:31:15 00:04:02 10.167.0.62

```


Only 1 receiver is displayed for a particular group, this is the device that responded to the IGMP query request from the FGT-3. If a receiver is active, the expire time should drop to approximately 2 minutes before being refreshed.

Checking the PIM-SM neighbors

Next, the PIM-SM neighbors should be checked. A PIM router becomes a neighbor when the PIM router receives a PIM hello. Use the following command to display the PIM-SM neighbors of FGT-3:

```
FGT-3 # get router info multicast pim sparse-mode neighbour
Neighbor Interface Uptime/Expires Ver DR
Address Priority/Mode
10.132.0.156 port2 01:57:12/00:01:33 v2 1 /
```

Checking that the PIM router can reach the RP

The rendezvous point (RP) must be reachable for the PIM router (FGT-3) to be able to send the *, G join to request the stream. This can be checked for FGT-3 using the following command:

```
FGT-3 # get router info multicast pim sparse-mode rp-mapping
PIM Group-to-RP Mappings
Group(s): 224.0.0.0/4, Static
RP: 192.168.1.1
Uptime: 07:23:00
```

Viewing the multicast routing table (FGT-3)

The FGT-3 unicast routing table can be used to determine the path taken to reach the RP at 192.168.1.1. You can then check the stream state entries using the following commands:

```
FGT-3 # get router info multicast pim sparse-mode table
IP Multicast Routing Table

(*,*,RP) Entries: 0
(*,G) Entries: 1
(S,G) Entries: 1
(S,G,rpt) Entries: 1
FCR Entries: 0
```

(*,*,RP) Entries	This state may be reached by general joins for all groups served by a specified RP.
(*,G) Entries	State that maintains the RP tree for a given group.
(S,G) Entries	State that maintains a source-specific tree for source S and group G.
(S,G,rpt) Entries	State that maintains source-specific information about source s on the RP tree for G. For example, if a source is being received on the source-specific tree, it will normally have been pruned off the RP tree.
FCR	The FCR state entries are for tracking the sources in the <*, G> when <S, G> isn't available for any reason, the stream would typically be flowing when this state exists.

Breaking down each entry in detail:

```
(*, 239.255.255.1)
RP: 192.168.1.1
RPF nbr: 10.132.0.156
RPF idx: port2
Upstream State: JOINED
Local:
port3
Joined:
Asserted:
FCR:
```

The RP will always be listed in a *, G entry, the RPF neighbor and interface index will also be shown. In this topology, these are the same in all downstream PIM routers. The state is active so the upstream state is joined.

In this case FGT-3 is the last hop router so the IGMP join is received locally on port3. There is no PIM outgoing interface listed for this entry as it is used for the upstream PIM join.

```
(10.166.0.11, 239.255.255.1)
RPF nbr: 10.132.0.156
RPF idx: port2
SPT bit: 1
Upstream State: JOINED
Local:
Joined:
Asserted:
Outgoing:
port3
```

This is the entry for the SPT, no RP IS listed. The S, G stream will be forwarded out of the stated outgoing interface.

```
(10.166.0.11, 239.255.255.1, rpt)
RP: 192.168.1.1
RPF nbr: 10.132.0.156
RPF idx: port2
Upstream State: NOT PRUNED
Local:
Pruned:
Outgoing:
```

The above S, G, RPT state is created for all streams that have both a S, G and a *, G entry on the router. This isn't pruned, in this case, because of the topology: the RP and source are reachable over the same interface.

Although not seen in this scenario, assert states may be seen when multiple PIM routers exist on the same LAN, which can lead to more than one upstream router having a valid forwarding state. Assert messages are used to elect a single forwarder from the upstream devices.

Viewing the PIM next-hop table

The PIM next-hop table is also very useful for checking the various states, it can be used to quickly identify the states of multiple multicast streams.

```
FGT-3 # get router info multicast pim sparse-mode next-hop
Flags: N = New, R = RP, S = Source, U = Unreachable
Destination Type Nexthop Nexthop Metric Pref Refcnt
Num Addr Ifindex
-----
10.166.0.11 ..S. 1 10.132.0.156 9 21 110 3
192.168.1.1 .R.. 1 10.132.0.156 9 111 110 2
```

Viewing the PIM multicast forwarding table

Also, you can check the multicast forwarding table showing the ingress and egress ports of the multicast stream.

```
FGT-3 # get router info multicast table

IP Multicast Routing Table
Flags: I - Immediate Stat, T - Timed Stat, F - Forwarder installed
Timers: Uptime/Stat Expiry
Interface State: Interface (TTL threshold)

(10.166.0.11, 239.255.255.1), uptime 04:02:55, stat expires 00:02:25
Owner PIM-SM, Flags: TF
Incoming interface: port2
Outgoing interface list:
port3 (TTL threshold 1)
```

Viewing the kernel forwarding table

Also, the kernel forwarding table can be verified, however this should give similar information to the above command:

```
FGT-3 # diag ip multicast mroute
grp=239.255.255.1 src=10.166.0.11 intf=9 flags=(0x10000000)[ ] status=resolved
last_assert=2615136 bytes=1192116 pkt=14538 wrong_if=0 num_ifs=1
index(ttl)=[6(1),]
```

Viewing the multicast routing table (FGT-2)

If you check the output on FGT-2, there are some small differences:

```
FGT-2 # get router info multicast pim sparse-mode table
IP Multicast Routing Table

(*,*,RP) Entries: 0
(*,G) Entries: 1
(S,G) Entries: 1
(S,G,rpt) Entries: 1
FCR Entries: 0

(*, 239.255.255.1)
RP: 192.168.1.1
RPF nbr: 0.0.0.0
RPF idx: None
Upstream State: JOINED
Local:
Joined:
external
Asserted:
FCR:
```

The *, G entry now has a joined interface rather than local because it received a PIM join from FGT-3 rather than a local IGMP join.

```
(10.166.0.11, 239.255.255.1)
RPF nbr: 10.130.0.237
RPF idx: internal
SPT bit: 1
Upstream State: JOINED
```

```

Local:
Joined:
external
Asserted:
Outgoing:
external

```

The `S, G` entry shows that we have received a join on the external interface and the stream is being forwarded out of this interface.

```

(10.166.0.11, 239.255.255.1, rpt)
RP: 192.168.1.1
RPF nbr: 0.0.0.0
RPF idx: None
Upstream State: PRUNED
Local:
Pruned:
Outgoing:
External

```

The `S, G, RPT` is different from FGT-3 because FGT-2 is the RP, it has pruned back the SPT for the RP to the first hop router.

Viewing the multicast routing table (FGT-1)

FGT-1 again has some differences with regard to the PIM-SM states. There's no `*, G` entry because it isn't in the path of a receiver and the RP.

```

FGT-1_master # get router info multicast pim sparse-mode table
IP Multicast Routing Table

(*,*,RP) Entries: 0
(*,G) Entries: 0
(S,G) Entries: 1
(S,G,rpt) Entries: 1
FCR Entries: 0

```

Below the `S, G` is the SPT termination because this FortiGate is the first hop router. The RPF neighbor always shows as `0.0.0.0` because the source is local to this device. Both the joined and outgoing fields show as external because the PIM join and the stream is egressing on this interface.

```

(10.166.0.11, 239.255.255.1)
RPF nbr: 0.0.0.0
RPF idx: None
SPT bit: 1
Upstream State: JOINED
Local:
Joined:
external
Asserted:
Outgoing:
external

```

The stream has been pruned back from the RP because the end-to-end SPT is flowing. In this case, there's no requirement for the stream to be sent to the RP.

```

(10.166.0.11, 239.255.255.1, rpt)
RP: 0.0.0.0
RPF nbr: 10.130.0.156
RPF idx: external

```

```

Upstream State: RPT NOT JOINED
Local:
Pruned:
Outgoing:

```

Example multicast DNAT configuration

The example topology shown and described below shows how to configure destination NAT (DNAT) for two multicast streams. Both of these streams originate from the same source IP address, which is 10.166.0.11. The example configuration keeps the streams separate by creating 2 multicast NAT policies.

In this example, the FortiGate devices have the following roles:

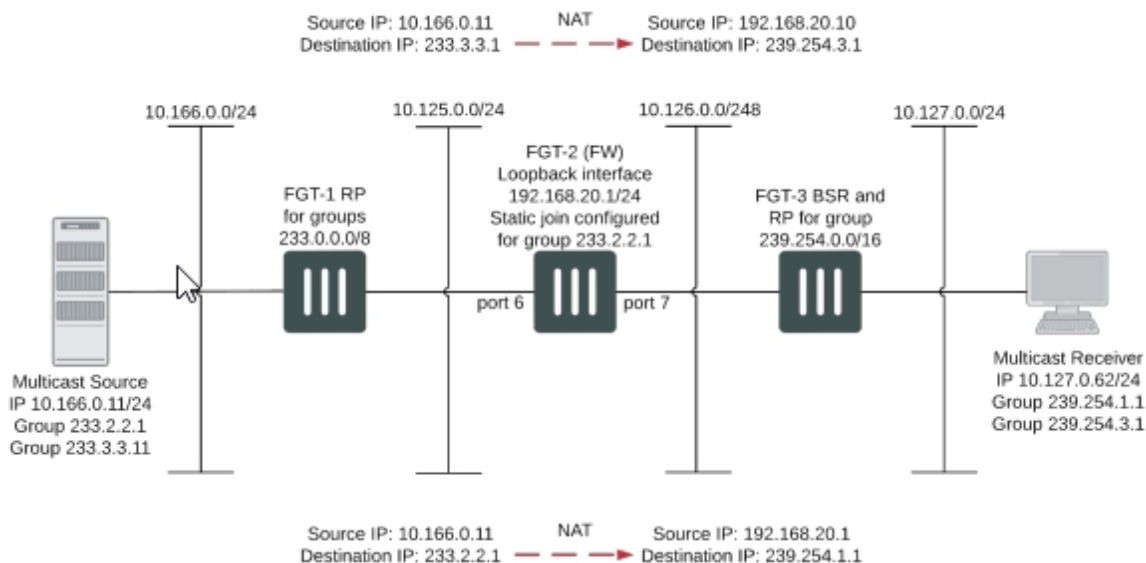
- FGT-1 is the RP for dirty networks, 233.0.0.0/8.
- FGT-2 performs all firewall and DNAT translations.
- FGT-3 is the RP for the clean networks, 239.254.0.0/16.
- FGT-1 and FGT-3 are functioning as PM enabled routers and could be replaced can be any PIM enabled router.

This example only describes the configuration of FGT-2.

FGT-2 performs NAT so that the receivers connected to FGT-3 receive the following translated multicast streams.

- If the multicast source sends multicast packets with a source and destination IP of 10.166.0.11 and 233.2.2.1, FGT-3 translates the source and destination IPs to 192.168.20.1 and 239.254.1.1
- If the multicast source sends multicast packets with a source and destination IP of 10.166.0.11 and 233.3.3.1, FGT-3 translates the source and destination IPs to 192.168.20.10 and 239.254.3.1

Example multicast DNAT topology



To configure FGT-2 for DNAT multicast

1. Add a loopback interface. In the example, the loopback interface is named `loopback`.

```
config system interface
```

```

edit loopback
    set vdom root
    set ip 192.168.20.1 255.255.255.0
    set type loopback
next
end

```

2. Add PIM and add a unicast routing protocol to the loopback interface as if it was a normal routed interface. Also, add static joins to the loopback interface for any groups to be translated.

```

config router multicast
config interface
edit loopback
    set pim-mode sparse-mode
    config join-group
        edit 233.2.2.1
        next
        edit 233.3.3.1
        next
    end
next

```

3. In this example, to add firewall multicast policies, different source IP addresses are required so you must first add an IP pool:

```

config firewall ippool
edit Multicast_source
    set endip 192.168.20.20
    set interface port6
    set startip 192.168.20.10
next
end

```

4. Add the translation security policies.

Policy 2, which is the source NAT policy, uses the actual IP address of port6. Policy 1, the DNAT policy, uses an address from the IP pool. The source and destination addresses will need to be previously created address objects. For this example, 233.3.3.1 255.255.255.255 will be represented by "example-addr_1" and 10.166.0.11 255.255.255.255 will be represented by "example-addr_2". You'll likely want to use something more intuitive from your own network.

```

config firewall multicast-policy
edit 1
    set dnat 239.254.3.1
    set dstaddr example-addr_1
    set dstintf loopback
    set nat 192.168.20.10
    set srcaddr example-addr_2
    set srcintf port6
next
edit 2
    set dnat 239.254.1.1
    set dstaddr 233.2.2.1 255.255.255.255
    set dstintf loopback
    set nat 192.168.20.1
    set srcaddr 10.166.0.11 255.255.255.255
    set srcintf port6
next
end

```

5. Add a firewall multicast policy to forward the stream from the loopback interface to the physical outbound interface.

This example is an any/any policy that makes sure traffic accepted by the other multicast policies can exit the FortiGate.

```
config firewall multicast-policy
  edit 3
    set dstintf port7
    set srcintf loopback
  next
end
```

Example PIM configuration that uses BSR to find the RP

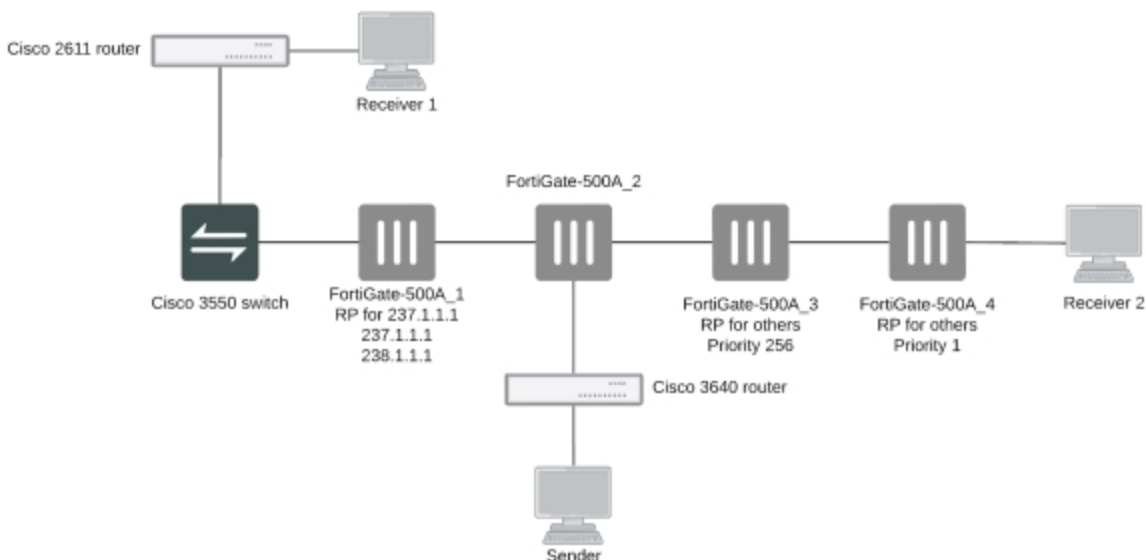
This example shows how to configure a multicast routing network for a network consisting of four FortiGate-500A devices (FortiGate-500A_1 to FortiGate-500A_4). A multicast sender is connected to FortiGate-500A_2. FortiGate-500A_2 forwards multicast packets in two directions to reach Receiver 1 and Receiver 2.

The configuration uses a Boot Start Router (BSR) to find the Rendezvous Points (RPs) instead of using static RPs. Under interface configuration, the loopback interface `lo0` must join the 236.1.1.1 group (source).

This example describes:

- Commands used in this example
- Configuration steps
- Example debug commands

PIM network topology using BSR to find the RP



Commands used in this example

This example uses CLI commands for the following configuration settings:

- Adding a loopback interface (lo0)
- Defining the multicast routing
- Adding the NAT multicast policy

Adding a loopback interface

Where required, the following command is used to define a loopback interface named lo0.

```
config system interface
  edit lo0
    set vdom root
    set ip 1.4.50.4 255.255.255.255
    set allowaccess ping https ssh snmp http telnet
    set type loopback
  next
end
```

Defining the multicast routing

In this example, the following command syntax is used to define multicast routing.

The example uses a Boot Start Router (BSR) to find the Rendezvous Points (RPs) instead of using static RPs. Under interface configuration, the loopback interface lo0 must join the 236.1.1.1 group (source).

```
config router multicast
  config interface
    edit port6
      set pim-mode sparse-mode
    next
    edit port1
      set pim-mode sparse-mode
    next
    edit lo0
      set pim-mode sparse-mode
      set rp-candidate enable
      config join-group
        edit 236.1.1.1
        next
      end
      set rp-candidate-priority 1
    next
  end
  set multicast-routing enable
  config pim-sm-global
    set bsr-allow-quick-refresh enable
    set bsr-candidate enable
    set bsr-interface lo0
    set bsr-priority 200
  end
end
```

Adding the NAT multicast policy

In this example, the incoming multicast policy does the address translation.

The NAT address should be the same as the IP address of the of loopback interface. The DNAT address is the translated address, which should be a new group.

```
config firewall multicast-policy
edit 1
set dstintf port6
set srcintf lo0
next
edit 2
set dnat 238.1.1.1
set dstintf lo0
set nat 1.4.50.4
set srcintf port1
next
```

Configuration steps

In this sample, FortiGate-500A_1 is the RP for the group 228.1.1.1, 237.1.1.1, 238.1.1.1, and FortiGate-500A_4 is the RP for the other group which has a priority of 1. OSPF is used in this example to distribute routes including the loopback interface. All firewalls have full mesh security policies to allow any to any.

- In the FortiGate-500A_1 configuration, the NAT policy translates source address 236.1.1.1 to 237.1.1.1
- In the FortiGate-500A_4 configuration, the NAT policy translates source 236.1.1.1 to 238.1.1.1
- Source 236.1.1.1 is injected into network as well

The following procedures include the CLI commands for configuring each of the FortiGate devices in the example configuration.

To configure FortiGate-500A_1

1. Configure multicast routing:

```
config router multicast
config interface
edit port5
set pim-mode sparse-mode
next
edit port4
set pim-mode sparse-mode
next
edit lan
set pim-mode sparse-mode
next
edit port1
set pim-mode sparse-mode
next
edit lo999
set pim-mode sparse-mode
next
edit lo0
set pim-mode sparse-mode
set rp-candidate enable
set rp-candidate-group 1
next
end
set multicast-routing enable
config pim-sm-global
```

```

        set bsr-candidate enable
        set bsr-interface lo0
    end
end

```

2. Add multicast security policies:

```

config firewall multicast-policy
    edit 1
        set dstintf port5
        set srcintf port4
    next
    edit 2
        set dstintf port4
        set srcintf port5
    next
    edit 3
    next
end

```

3. Add router access lists:

```

config router access-list
    edit 1
        config rule
            edit 1
                set prefix 228.1.1.1 255.255.255.255
                set exact-match enable
            next
            edit 2
                set prefix 237.1.1.1 255.255.255.255
                set exact-match enable
            next
            edit 3
                set prefix 238.1.1.1 255.255.255.255
                set exact-match enable
            next
        end
    next
end

```

To configure FortiGate-500A_2

1. Configure multicast routing:

```

config router multicast
    config interface
        edit "lan"
            set pim-mode sparse-mode
        next
    edit "port5"
        set pim-mode sparse-mode
    next
    edit "port2"
        set pim-mode sparse-mode
    next
    edit "port4"
        set pim-mode sparse-mode
    next
end

```

```
edit "lo_5"
    set pim-mode sparse-mode
    config join-group
        edit 236.1.1.1
        next
    end
next
end
set multicast-routing enable
end
```

2. Add multicast security policies:

```
config firewall multicast-policy
edit 1
    set dstintf lan
    set srcintf port5
next
edit 2
    set dstintf port5
    set srcintf lan
next
edit 4
    set dstintf lan
    set srcintf port2
next
edit 5
    set dstintf port2
    set srcintf lan
next
edit 7
    set dstintf port1
    set srcintf port2
next
edit 8
    set dstintf port2
    set srcintf port1
next
edit 9
    set dstintf port5
    set srcintf port2
next
edit 10
    set dstintf port2
    set srcintf port5
next
edit 11
    set dnat 237.1.1.1
    set dstintf lo_5
    set nat 5.5.5.5
    set srcintf port2
next
edit 12
    set dstintf lan
    set srcintf lo_5
next
edit 13
    set dstintf port1
```

```
        set srcintf lo_5
    next
    edit 14
        set dstintf port5
        set srcintf lo_5
    next
    edit 15
        set dstintf port2
        set srcintf lo_5
    next
    edit 16
    next
end
```

To configure FortiGate-500A_3

1. Configure multicast routing:

```
config router multicast
    config interface
        edit port5
            set pim-mode sparse-mode
        next
        edit port6
            set pim-mode sparse-mode
        next
        edit lo0
            set pim-mode sparse-mode
            set rp-candidate enable
            set rp-candidate-priority 255
        next
        edit lan
            set pim-mode sparse-mode
        next
    end
set multicast-routing enable
config pim-sm-global
    set bsr-candidate enable
    set bsr-interface lo0
end
end
```

2. Add multicast security policies:

```
config firewall multicast-policy
    edit 1
        set dstintf port5
        set srcintf port6
    next
    edit 2
        set dstintf port6
        set srcintf port5
    next
    edit 3
        set dstintf port6
        set srcintf lan
    next
    edit 4
        set dstintf lan
```

```
        set srcintf port6
    next
    edit 5
        set dstintf port5
        set srcintf lan
    next
    edit 6
        set dstintf lan
        set srcintf port5
    next
end
```

To configure FortiGate-500A_4

1. Configure multicast routing:

```
config router multicast
config interface
    edit port6
        set pim-mode sparse-mode
    next
    edit lan
        set pim-mode sparse-mode
    next
    edit port1
        set pim-mode sparse-mode
    next
    edit lo0
        set pim-mode sparse-mode
        set rp-candidate enable
        config join-group
            edit 236.1.1.1
            next
        end
        set rp-candidate-priority 1
    next
end
set multicast-routing enable
config pim-sm-global
set bsr-allow-quick-refresh enable
set bsr-candidate enable
set bsr-interface lo0
set bsr-priority 1
end
end
```

2. Add multicast security policies:

```
config firewall policy
    edit 1
        set srcintf lan
        set dstintf port6
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule always
        set service ANY
    next
    edit 2
```

```
        set srcintf port6
        set dstintf lan
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule always
        set service ANY
    next
    edit 3
        set srcintf port1
        set dstintf port6
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule always
        set service ANY
    next
    edit 4
        set srcintf port6
        set dstintf port1
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule always
        set service ANY
    next
    edit 5
        set srcintf port1
        set dstintf lan
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule always
        set service ANY
    next
    edit 6
        set srcintf lan
        set dstintf port1
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule always
        set service ANY
    next
    edit 7
        set srcintf port1
        set dstintf port1
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule always
        set service ANY
    next
    edit 8
        set srcintf port6
        set dstintf lo0
        set srcaddr all
```

```
        set dstaddr all
        set action accept
        set schedule always
        set service ANY
    next
    edit 9
        set srcintf port1
        set dstintf lo0
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule always
        set service ANY
    next
    edit 10
        set srcintf lan
        set dstintf lo0
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule always
        set service ANY
    next
end
```

SD-WAN

This chapter describes FortiOS 6.0 features for SD-WAN.

Application control settings in SD-WAN rules

You can configure application control and application control groups in SD-WAN rules.

To configure application control in SD-WAN rules - GUI:

1. Go to **Network > SD-WAN Rules**.
2. Select **Create New** or **Edit** an existing rule.
3. In the **Destination** section, in **Destination type** field, select **Internet Service**.
4. In the **Select Entries** side menu, select **Application Control**.
5. Select one or more applications from the list.

To configure application control in SD-WAN rules - CLI:

```
config application group
edit <name>
    set application <application-ID-list>
next
end
```

Internet services in SD-WAN rules

Internet service support is expanded from a single Internet service to support Internet service groups, custom Internet service groups, control-based application identification, and control-based Internet service groups.

To configure Internet services in SD-WAN rules - GUI:

1. Go to **Network > SD-WAN Rules**.
2. **Create New** or **Edit** an existing rule.
3. In the **Destination** section, in **Destination type** field, select **Internet Service**.
4. In the **Select Entries** side menu, select **Internet Service**.
5. Select one or more Internet services from the list.

To configure Internet service groups in SD-WAN rules - CLI:

```
config system virtual-wan-link
config service
edit <priority-rule-ID>
    set internet-service enable
    set internet-service-group <Internet-service-group-list>
    set internet-service-custom-group <custom-Internet-service-group-list>
    set internet-service-ctrl <control-based-Internet-service-ID-list>
    set internet-service-ctrl-group <control-based-Internet-service-group-list>
next
end
```



```
end
```

To configure Internet service groups - CLI:

```
config firewall internet-service-group
  edit <Internet-service-group-name>
    set comment [string]
    set member <Internet-service-group-members>
  next
end
```

To configure custom Internet service groups - CLI:

```
config firewall internet-service-custom-group
  edit <custom-Internet-service-group-name>
    set comment [string]
    set member <custom-Internet-service-group-members>
  next
end
```

Bandwidth and custom profile options in SD-WAN rules

You can configure bandwidth options for the link cost factor, so that a FortiGate selects the link based on available bandwidth of incoming, outgoing, or bidirectional traffic. This is useful because users may use some applications primarily for downloading and other applications primarily for uploading.

You can also configure a custom profile for the link cost factor. A FortiGate selects the best link using the following formula: $\text{Link quality} = (\text{packet-loss-weight} * \text{packet loss}) + (\text{latency-weight} * \text{latency}) + (\text{jitter-weight} * \text{jitter}) + (\text{bandwidth-weight} / \text{bandwidth})$.

To configure bandwidth and custom profile options in SD-WAN rules - GUI:

1. Go to **Network > SD-WAN Rules**.
2. In the **Outgoing Interfaces** section, in the **Strategy** field, select **Best Quality**.
3. After you add interfaces in the **Interface preference** field, a **Quality criteria** field appears.
4. In the **Quality criteria** field, you can select **Downstream** for downstream bandwidth, **Upstream** for upstream bandwidth, **Bandwidth** for bidirectional bandwidth, or **custom-profile-1** to create a custom profile.
5. After you select **custom-profile-1**, you can set the **latency-weight**, **jitter-weight**, **packet-loss-weight**, and **bandwidth-weight**.

To configure bandwidth options in SD-WAN rules - CLI:

```
config system virtual-wan-link
  config service
    edit <priority-rule-ID>
      set link-cost-factor inboundwidth
      set link-cost-factor outboundwidth
      set link-cost-factor bibandwidth
    next
  end
end
```

To configure a custom profile in SD-WAN rules - CLI:

```

config system virtual-wan-link
config service
    edit <priority-rule-ID>
        set link-cost-factor custom-profile-1
        set packet-loss-weight <integer>
        set latency-weight <integer>
        set jitter-weight <integer>
        set bandwidth-weight <integer>
    next
end
end

```

where packet-loss-weight, latency-weight, jitter-weight, and bandwidth-weight are integers in the range of 0 to 10000000.

SLA management

You can configure service level agreement (SLA) management in the FortiGate GUI.

If all links meet the SLA criteria, the FortiGate uses the first link, even if that link isn't the best quality link. If at any time, the link in use doesn't meet the SLA criteria, and the next link in the configuration meets the SLA criteria, the FortiGate changes to that link. If the next link doesn't meet the SLA criteria, the FortiGate uses the next link in the configuration if it meets the SLA criteria, and so on.

If no links meet the SLA criteria, the FortiGate uses the preferred link, which is the first link in the configuration. FortiGate continually checks the links for to see if they meet the SLA criteria.

To configure SLA - GUI:

1. Go to **Network > Performance SLA**.
2. Select **Create New**.
3. In the **Name**, **Protocol**, **Server**, and **Participants** fields, set appropriate values.
4. In the **SLA Targets** section, set appropriate values in the **Latency threshold**, **Jitter threshold**, and **Packet loss threshold** fields. Select **OK**.
The Performance SLA page changes to display information about the SLA.
5. Go to **Network > SD-WAN Rules**.
6. In the **Name** field, set a name for the rule.
7. In the **Destination type** field, select **Internet Service**. In the **Destination** field, select **+**. In the **Select Entries** window, select **Internet Service** and select the appropriate Internet services from the list. Select **Close**.
8. In the **Outgoing Interfaces** section, in the **Strategy** field, select **Minimum Quality (SLA)**.
9. In the **Interface preference** field, select **+**. In the **Select Entries** window, select the appropriate interfaces from the list. Select **Close**.
10. In the **Required SLA target** field, select the appropriate SLA from the drop-down list. Select **OK**.
The SD-WAN Rules page changes to display information about the SD-WAN Rules. You can drag and drop the rules to reorder them.

To configure SLA in health checks - CLI:

```

config system virtual-wan-link

```

```

config health-check
  edit <health-check-name>
    config sla
      edit <SLA-ID>
        set link-cost-factor {latency | jitter | packet-loss}
        set latency-threshold <milliseconds>
        set jitter-threshold <milliseconds>
        set packetloss-threshold <percentage>
      next
    end
  end
end

```

To configure SLA in services - CLI:

```

config system virtual-wan-link
  config service
    edit <priority-rule-ID>
      set mode sla
      config sla
        edit <health-check-name>
          set id <SLA-ID>
        next
      end
    next
  end
end

```

To set link priority mode to SLA in SD-WAN rules - CLI:

```

config system virtual-wan-link
  config service
    edit <priority-rule-ID>
      set mode sla
      set link-cost-threshold <threshold-change-percentage>
      set priority-members <member-sequence-list>
    next
  end
end

```

Multiple server support for health checks

You can now configure multiple servers for health checks in SD-WAN.

To configure server support for health checks - CLI:

```

config system virtual-wan-link
  config health-check
    edit <health-check-name>
      set server <server-list>
    next
  end
end

```

where <server-list> is one or more IP addresses or FQDN names of servers.

IPv6 support for SD-WAN

SD-WAN now supports IPv6. It supports all load balance modes, health checking (ping6), and service rules for source address, source user and group, and destination address.

FortiOS 6.0 also increases the configuration limit for health checks and priority rules in SD-WAN. The limit for both health checks and priority rules is increased from 256 to 4096 globally and 512 to 4096 per VDOM.

To configure SD-WAN member interfaces - CLI:

```
config system virtual-wan-link
  config members
    edit <sequence-number>
      set interface <interface-name>
      set {gateway | gateway6} <gateway-address>
    next
  end
end
```

To enable SD-WAN - CLI:

```
config router {static | static6}
  edit <sequence-number>
    set virtual-wan-link enable
  next
end
```

To configure health check (IPv4) - CLI:

```
config system virtual-wan-link
  config health-check
    edit <health-check-name>
      set address-mode ipv4
      set protocol {ping | tcp-echo | udp-echo | http | twamp}
    next
  end
end
```

To configure health check (IPv6) - CLI:

```
config system virtual-wan-link
  config health-check
    edit <health-check-name>
      set address-mode ipv6
      set protocol ping6
    next
  end
end
```

To configure service rules - CLI:

```
config system virtual-wan-link
  config service
    edit <priority-rule-ID>
      set name <rule-name>
```

```

        set addr-mode {ipv4 | ipv6}
        set member <sequence-number>
        set {dst | dst6} <destination-address-name>
        set {src | src6} <source-address-name>
    next
end
end

```

DSCP tagging of forwarded packets in SD-WAN rules

You can now configure DSCP tagging of forwarded packets based on identified applications in SD-WAN rules.

To configure DSCP tagging of forwarded packets in SD-WAN rules - CLI:

```

config system virtual-wan-link
config service
    edit <priority-rule-ID>
        set dscp-forward {enable | disable}
        set dscp-reverse {enable | disable}
        set dscp-forward-tag <binary>
        set dscp-reverse tag <binary> [6 bits binary, range 000000-111111]
    next
end
end

```

where:

- `dscp-forward-tag <binary>` is the forward traffic DSCP tag. It's a 6 bit binary value in the range of 000000-111111.
- `dscp-reverse tag <binary>` is the reverse traffic DSCP tag. It's a 6 bit binary value in the range of 000000-111111.

SD-WAN and dynamic routing

Dynamic routing is now supported for SD-WAN. You can use dynamic routing and SD-WAN. Using BGP, SD-WAN can update its rules using dynamic routes.

To set a route tag in route map rules - CLI:

```

config router route-map
    edit <name>
        config rule
            edit <rule-ID>
                set match-community <BGP-community-list>
                set set-route-tag <route-tag>
            next
        end
    next
end

```

where `<route-tag>` is a number from 0 to 4294967295.

To set health check members - CLI:

```

config system virtual-wan-link
config health-check
    edit <health-check-name>

```

```

        set members <member-sequence-number-list>
    next
end
end

```

To add a route tag in SD-WAN rules - CLI:

```

config system virtual-wan-link
config service
    edit <priority-rule-ID>
        set route-tag <route-map-route-tag>
    next
end
end

```

where <route-tag> is a number from 0 to 4294967295.

Link priority in SD-WAN rules

You can configure the priority for links in SD-WAN rules. You can use the `set mode` command to control how the priority rule sets the priority of interfaces in SD-WAN. If you set this to `priority`, a FortiGate assigns priorities to the interfaces based on the order of the interfaces that you configure using the `set priority-members` command.

In priority mode, when the difference between two links is within the amount that you configure for the `link-cost-threshold`, the FortiGate uses the link with the higher priority, which is the first member in the `priority-members` list.

In SLA mode, the FortiGate assigns the link based on SLA settings.

To set link priority mode in SD-WAN rules - CLI:

```

config system virtual-wan-link
config service
    edit <priority-rule-ID>
        set mode priority
        set link-cost-threshold <threshold-change-percentage>
        set priority-members <member-sequence-list>
    next
end
end

```

SD-WAN rule for address negation

When traffic travels through a FortiGate, it can match a software-defined wide area network (SD-WAN) rule first, even if it's meant for another interface, resulting in a traffic interruption. To avoid this, you can configure an implicit rule for address negation so that SD-WAN policy-based routing (PBR) rules don't match traffic unless the traffic is meant for SD-WAN interfaces.

You can enable address negation for both destination and source address matches. During the process of matching rules, if the rule is an SD-WAN rule, the FortiGate checks the outgoing interface first. If the interface isn't an SD-WAN member interface, the rule is ignored.

To enable the negation of destination address match - CLI:

```

config system virtual-wan-link

```

```
config service
  edit <id>
    set dst-negate enable
  next
end
```

To enable the negation of source address match - CLI:

```
config system virtual-wan-link
  config service
    edit <id>
      set src-negate enable
    next
end
```

SD-WAN CLI changes

The following CLI changes have been made for SD-WAN:

- Removed the timeout option from the link health monitor. You can no longer use the `set timeout` command under the `config system link-monitor` command.
- Removed the timeout option from the health check. You can no longer use the `set timeout` command under the `config health-check` command.

Troubleshooting

Netflow support

Netflow is a networking feature, introduced by Cisco, to collect and export information about traffic flow through routers. IPFIX (Internet Protocol Flow Information Export) is the standardized Internet Protocol based on NetFlow version 9. The standards requirements for IPFIX are outlined in RFC 3197, and its basic specifications and other information are documented in RFC 5103, RFC 6759, and RFC 7011 through RFC 7015.

To enable and configure NetFlow traffic - CLI:

```
config system netflow
  set collector-ip <collector IP>
  set collector-port <NetFlow collector port>
  set csource-ip <Source IP for NetFlow agent>
  set cactive-flow-timeout <time in minutes of timeout to report active flows>
  set cinactive-flow-timeout <time in seconds of timeout for periodic report of finished flows>
end
```

You can also configure these settings per VDOM, using the following CLI command:

```
config system vdom-netflow
```

You must also enable a Netflow sampler on specific interfaces.

sFlow support

sFlow is a method of monitoring the traffic on your network to identify areas on the network that may impact performance and throughput. FortiOS implements sFlow version 5.

sFlow uses packet sampling to monitor network traffic. The sFlow Agent captures packet information at defined intervals and sends them to an sFlow Collector for analysis, providing real-time data analysis. The information sent is only a sampling of the data for minimal impact on network throughput and performance.

The sFlow Agent is embedded in the FortiGate. Once configured, the FortiGate sends sFlow datagrams of the sampled traffic to the sFlow Collector, also called an sFlow Analyzer. The sFlow Collector receives the datagrams, and provides real-time analysis and graphing to indicate where potential traffic issues are occurring. sFlow Collector software is available from a number of third party software vendors.

sFlow data captures only a sampling of network traffic, not all traffic like the traffic logs on the FortiGate. Sampling works by the sFlow Agent looking at traffic packets when they arrive on an interface. A decision is made whether the packet is dropped and allowed to be to its destination or if a copy is forwarded to the sFlow Collector. The sample used and its frequency are determined during configuration.

sFlow is not supported on virtual interfaces such as vdom link, ipsec, ssl.root or gre.

The sFlow datagram sent to the Collector contains the information:

- Packet header (e.g. MAC,IPv4,IPv6,IPX,AppleTalk,TCP,UDP, ICMP)
- Sample process parameters (rate, pool, etc.)
- Input/output ports

- Priority (802.1p and TOS)
- VLAN (802.1Q)
- Source/destination prefix
- Next hop address
- Source AS, Source Peer AS
- Destination AS Path
- Communities, local preference
- User IDs (TACACS/RADIUS) for source/destination
- URL associated with source/destination
- Interface statistics (RFC 1573, RFC 2233, and RFC 2358)

sFlow agents can be added to any type of FortiGate interface. sFlow isn't supported on some virtual interfaces such as VDOM link, IPsec, gre, and ssl.root.

For more information on sFlow, Collector software and sFlow MIBs, visit www.sflow.org.

Configuration

sFlow configuration is available only in the CLI. Configuration requires two steps: enabling the sFlow Agent and configuring the interface for the sampling information.

To enable sFlow - CLI:

```
config system sflow
    set collector-ip <ip_address>
    set collector-port <port_number>
    set source-ip <ip_address>
end
```

The default port for sFlow is UDP 6343.

To configure in VDOM - CLI:

```
config system vdom-sflow
    set vdom-sflow enable
    set collector-ip <ip_address>
    set collector-port <port_number>
    set source-ip <ip_address>
end
```

To configure sFlow agents per interface - CLI:

```
config system interface
    edit <interface_name>
        set sflow-sampler enable
        set sample-rate <every_n_packets>
        set sample-direction [tx | rx | both]
        set polling-interval <seconds>
    next
end
```

Packet capture

When troubleshooting networks, it helps to look inside the header of the packets. This helps to determine if the packets, route, and destination are all what you expect. Packet capture can also be called a network tap, packet sniffing, or logic analyzing.

To use the packet capture - GUI:

1. Go to **Network > Packet Capture**.
2. Select **Create New** or select an existing entry if you have already made one that fits your needs.
3. Select the interface to monitor and select the number of packets to keep.
4. Select **Enable Filters**.
5. Enter the information you want to gather from the packet capture.
6. Select **OK**.

To run the capture, select the play button in the progress column in the packet capture list. If it's not active, **Not Running** also appears in the column cell. The progress bar indicates the status of the capture. You can stop and restart it at any time.

When the capture is complete, select the **Download** icon to save the packet capture file to your hard disk for further analysis.

Packet capture tells you what is happening on the network at a low level. This can be very useful for troubleshooting problems, such as:

- Finding missing traffic
- Seeing if sessions are setting up properly
- Locating ARP problems such as broadcast storm sources and causes
- Confirming which address a computer is using on the network if they have multiple addresses or are on multiple networks
- Confirming routing is working as you expect
- Wireless client connection problems
- Intermittent missing PING packets
- A particular type of packet is having problems, such as UDP, which is commonly used for streaming video

If you're running a constant traffic application, such as ping, packet capture can tell you if the traffic is reaching the destination, how the port enters and exits the FortiGate, if the ARP resolution is correct, and if the traffic is returning to the source as expected. You can also use packet switching to verify that NAT, or other configuration, is translating addresses or routing traffic the way that you want it to.

Before you start capturing packets, you need to have a good idea of what you are looking for. Capture is used to confirm or deny your ideas about what is happening on the network. If you try capture without a plan to narrow your search, you can end up with too much data to effectively analyze. On the other hand, you need to capture enough packets to really understand all of the patterns and behavior that you're looking for.

Chapter 18 - Parallel Path Processing (Life of a Packet)

This FortiOS Handbook chapter contains the following sections:

[Parallel Path Processing](#) introduces the concept of Parallel Path Processing.

[Packet flow ingress and egress: FortiGates without network processor offloading](#) describes the overall packet flow through a FortiGate with no network offloading (NP) hardware.

[Packet flow: FortiGates with NP6 processors first packet of a new session](#) similar to the previous section, the first packet in a new session that can be offloaded is processed in much the same way as on a FortiGate with no network processors.

[Packet flow: FortiGates with NP6 processors - packets in an offloaded session](#) describes the much simpler packet flow for a packet from an offloaded session.

[UTM/NGFW packet flow: flow-based inspection](#) describes how single pass UTM/NGFW processing occurs in a flow-based FortiGate or VDOM.

[UTM/NGFW packet flow: proxy-based inspection](#) describes how UTM/NGFW processing occurs in a proxy-based FortiGate or VDOM.

[UTM/NGFW packet flow: explicit web proxy](#) describes how Explicit web proxy processing occurs.

[Comparison of inspection types](#) shows how different security functions map to different inspection types.

Parallel Path Processing

Parallel Path Processing (PPP) uses the firewall policy configuration to choose from a group of parallel options to determine the optimal path for processing a packet. Most FortiOS features are applied through Firewall policies and the features applied determine the path a packet takes. Using firewall policies you can impose UTM/NGFW processing on content traffic that may contain security threats (such as HTTP, email and so on). Many UTM/NGFW processes are offloaded and accelerated by CP8 or CP9 processors. Using the policy configuration you can apply a range of protection from basic IPS attack protection that looks for network-based attacks to full scale advanced threat management (ATM), application control, antivirus, DLP and so on.

You can also create policies for traffic that does not pose security threats and bypass UTM/NGFW checking. This control allows you to improve network performance without compromising security. On FortiGates with network processors (for example the NP6) much of the traffic that does not require UTM/NGFW processing can be offloaded to the NP6 processors freeing up FortiGate processing resources for other higher risk traffic.

In addition, many FortiGate models support NTurbo to offload flow-based UTM/NGFW sessions to network processors. Flow-based sessions can also be accelerated using IPSA technology to enhance offloading of pattern matching to CP8 and CP9 content processors.

This chapter begins with an overview of packet flow ingress and egress and includes a section that shows how NP6 offloading optimizes packet flow for packets that don't require UTM/NGFW processing and for packets that use NTurbo to offload flow-based UTM/NGFW processing.

Next this chapter breaks down how packets pass through UTM/NGFW processing both for a single-pass flow-based UTM/NGFW processing and a proxy-based UTM/NGFW processing.

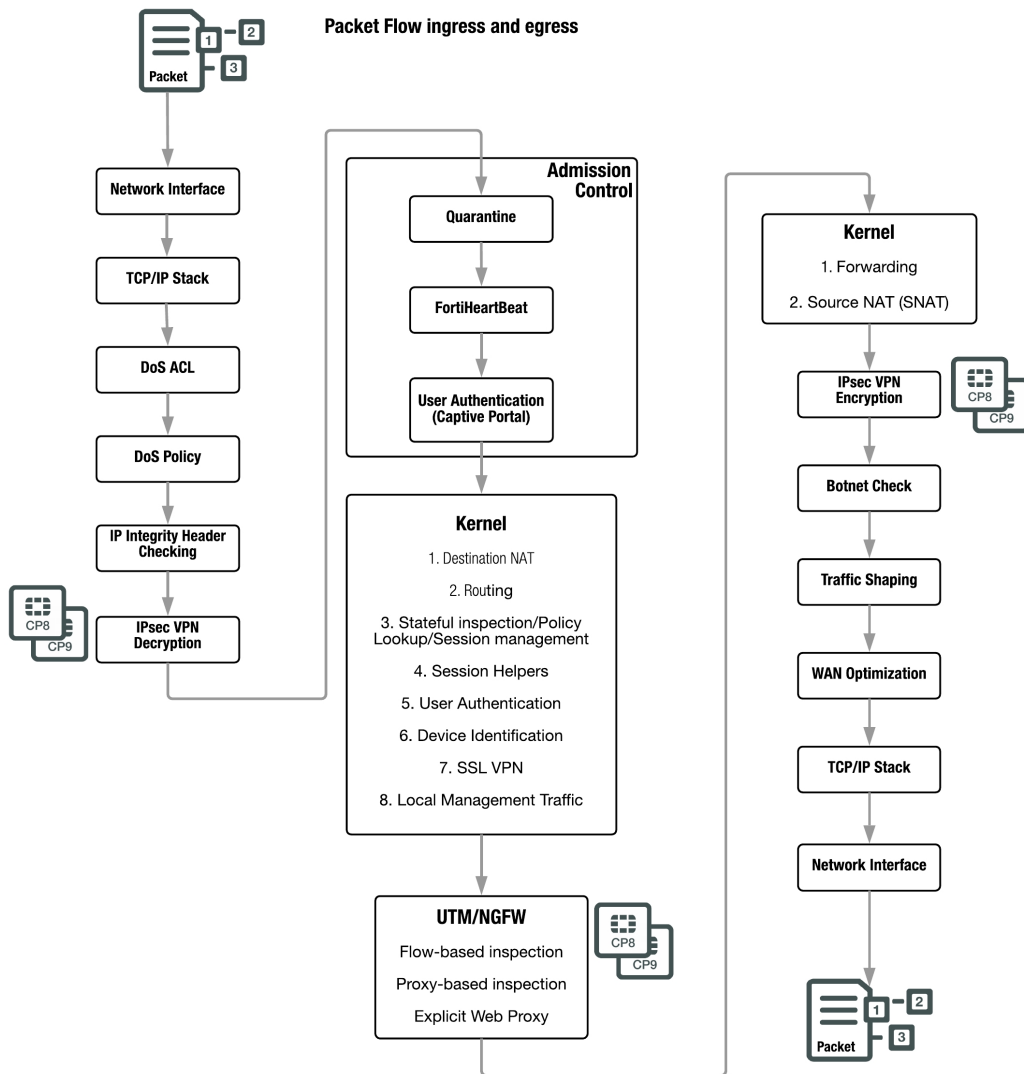
High-level list of processes that affect packets

In general packets passing through a FortiGate can be affected by the following processes. This is a complete high-level list of all of the processes. Not all packets see all of these processes. The processes a packet encounters depends on the type of packet and on the FortiGate software and hardware configuration.

- **Ingress packet flow**
 - Network Interface
 - TCP/IP stack
 - DoS ACL
 - DoS Policy
 - IP integrity header checking
 - IPsec VPN decryption
- **Admission Control**
 - Quarantine
 - FortiTelemetry
 - User Authentication
- **Kernel**
 - Destination NAT
 - Routing
 - Stateful inspection/Policy
Lookup/Session management
 - Session Helpers
 - User Authentication
 - Device Identification
 - SSL VPN
 - Local Management Traffic
- **UTM/NGFW**
 - Flow-based inspection
 - NTurbo
 - IPSA
 - Proxy-based inspection
 - Explicit Web Proxy
- **Kernel**
 - Forwarding
 - Source NAT (SNAT)
- **Egress packet flow**
 - IPsec VPN Encryption
 - Botnet check
 - Traffic shaping
 - WAN Optimization
 - TCP/IP stack
 - Network Interface

Packet flow ingress and egress: FortiGates without network processor offloading

This section describes the steps a packet goes through as it enters, passes through and exits from a FortiGate. This scenario shows all of the steps a packet goes through if a FortiGate does not contain network processors (such as the NP6).



Ingress

All packets accepted by a FortiGate pass through a network interface and are processed by the TCP/IP stack. Then if **DoS policies** or **Access Control List (ACL) policies** have been configured the packet must pass through these as well as automatic **IP integrity header checking**.

DoS scans are handled very early in the life of the packet to determine whether the traffic is valid or is part of a DoS attack. The DoS module inspects all traffic flows but only tracks packets that can be used for DoS attacks (for example, TCP SYN packets), to ensure they are within the permitted parameters. Suspected DoS attacks are blocked, other packets are allowed.

IP integrity header checking reads the packet headers to verify if the packet is a valid TCP, UDP, ICMP, SCTP or GRE packet. The only verification that is done at this step to ensure that the protocol header is the correct length. If it is, the packet is allowed to carry on to the next step. If not, the packet is dropped.

Incoming **IPsec packets** that match configured IPsec tunnels on the FortiGate are decrypted after header checking is done.

If the packet is an IPsec packet, the IPsec engine attempts to decrypt it. If the IPsec engine can apply the correct encryption keys and decrypt the packet, the unencrypted packet is sent to the next step. Non-IPsec traffic and IPsec traffic that cannot be decrypted passes on to the next step without being affected. IPsec VPN decryption is offloaded to and accelerated by CP8 or CP9 processors.

Admission control

Admission control checks to make sure the packet is not from a source or headed to a destination on the quarantine list. If configured admission control then imposes FortiTelemetry protection that requires a device to have FortiClient installed before allowing packets from it. Admission control can also impose captive portal authentication on ingress traffic.

Kernel

Once a packet makes it through all of the ingress steps, the FortiOS kernel performs the following checks to determine what happens to the packet next.

Destination NAT

Destination NAT checks the NAT table and determines if the destination IP address for incoming traffic must be changed using DNAT. DNAT is typically applied to traffic from the internet that is going to be directed to a server on a network behind the FortiGate. DNAT means the actual address of the internal network is hidden from the internet. This step determines whether a route to the destination address actually exists. DNAT must take place before routing so that the FortiGate can route packets to the correct destination.

Routing

Routing uses the routing table to determine the interface to be used by the packet as it leaves the FortiGate. Routing also distinguishes between local traffic and forwarded traffic. Firewall policies are matched with packets depending on the source and destination interface used by the packet. The source interface is known when the packet is received and the destination interface is determined by routing.

Stateful inspection/policy lookup/session management

Stateful inspection looks at the first packet of a session and looks in the policy table to make a security decision about the entire session. Stateful inspection looks at packet TCP SYN and FIN flags to identify the start and end of a session, the source/destination IP, source/destination port and protocol. Other checks are also performed on the packet payload and sequence numbers to verify it as a valid session and that the data is not corrupted or poorly formed.

When the first packet of a session is matched in the policy table, stateful inspection adds information about the session to its session table. So when subsequent packets are received for the same session, stateful inspection can determine how to handle them by looking them up in the session table (which is more efficient than looking them up in the policy table).

Stateful inspection makes the decision to drop or allow a session and apply security features to it based on what is found in the first packet of the session. Then all subsequent packets in the same session are processed in the same way.

When the final packet in the session is processed, the session is removed from the session table. Stateful inspection also has a session idle timeout that removes sessions from the session table that have been idle for the length of the timeout.

See the Stateful Firewall Wikipedia article (https://en.wikipedia.org/wiki/Stateful_firewall) for an excellent description of stateful inspection.

Session helpers

Some protocols include information in the packet body (or payload) that must be analyzed to successfully process sessions for this protocol. For example, the SIP VoIP protocol uses TCP control packets with a standard destination port to set up SIP calls. To successfully process SIP VoIP calls, FortiOS must be able to extract information from the body of the SIP packet and use this information to allow the voice-carrying packets through the firewall.

FortiOS uses session helpers to analyze the data in the packet bodies of some protocols and adjust the firewall to allow those protocols to send packets through the firewall. FortiOS includes the following session helpers:

- PPTP
- H323
- RAS
- TNS
- TFTP
- RTSP
- FTP
- MMS
- PMAP
- SIP
- DNS-UDP
- RSH
- DCERPC
- MGCP

User authentication

User authentication added to security policies is handled by the stateful inspection, which is why Firewall authentication is based on IP address. Authentication takes place after policy lookup selects a policy that includes authentication.

Device identification

Device identification is applied if required by the matching policy.

SSL VPN

Local SSL VPN traffic is treated like special management traffic as determined by the SSL VPN destination port. Packets are decrypted and are routed to an SSL VPN interface. Policy lookup is then used to control how packets are forwarded to their destination outside the FortiGate. SSL encryption and decryption is offloaded to and accelerated by CP8 or CP9 processors.

Local management traffic

Local management traffic terminates at a FortiGate interface. This can be any FortiGate interface including dedicated management interfaces. In multiple VDOM mode local management traffic terminates at the management interface. In transparent mode, local management traffic terminates at the management IP address.

Local management traffic includes administrative access, some routing protocol communication, central management from FortiManager, communication with the FortiGuard network and so on. Management traffic is allowed or blocked according to the Local In Policy list which lists all management protocols and their access control settings. You configure local management access indirectly by configuring administrative access and so on.

Management traffic is processed by applications such as the web server which displays the FortiOS GUI, the SSH server for the CLI or the FortiGuard server to handle local FortiGuard database updates or FortiGuard Web Filtering URL lookups.

Local management traffic is not involved in subsequent stateful inspection steps.

SSL VPN traffic terminates at a FortiGate interface similar to local management traffic. However, SSL VPN traffic uses a different destination port number than administrative HTTPS traffic and can thus be detected and handled differently.

UTM/NGFW

If the policy matching the packet includes security profiles, then the packet is subject to Unified Threat Management (UTM)/Next Generation Firewall (NGFW) processing. UTM/NGFW processing depends on the inspection mode of the FortiGate: Flow-based (single pass architecture) or proxy-based. Proxy-based processing can include Explicit web proxy traffic. Many UTM/NGFW processes are offloaded and accelerated by CP8 or CP9 processors.

Single pass flow-based UTM/NGFW inspection identifies and blocks security threats in real time as they are identified using single-pass Direct Filter Approach (DFA) pattern matching to identify possible attacks or threats.

Proxy-based UTM/NGFW inspection can apply both flow-based and proxy-based inspection. Packets initially encounter the IPS engine, which can apply single-pass flow-based IPS and Application Control (as configured). The packets are then sent to the proxy for proxy-based inspection. Proxy-based inspection can apply VoIP inspection, DLP, AntiSpam, Web Filtering, Antivirus, and ICAP.

Explicit web proxy inspection is similar to proxy based inspection.

Content processors (CP8 and CP9)

Most FortiGate models contain Security Processing Unit (SPU) Content Processors (CPs) that accelerate many common resource intensive security related processes. CPs work at the system level with tasks being offloaded to them as determined by the main CPU. Capabilities of the CPs vary by model. Newer FortiGate units include CP8 and CP9 processors. Older CP versions still in use in currently operating FortiGate models include the CP4, CP5, and CP6.

CP9 capabilities

The CP9 content processor provides the following services:

- Flow-based inspection (IPS, application control etc.) pattern matching acceleration with over 10Gbps throughput
 - IPS pre-scan
 - IPS signature correlation
 - Full match processors
- High performance VPN bulk data engine
 - IPsec and SSL/TLS protocol processor
 - DES/3DES/AES128/192/256 in accordance with FIPS46-3/FIPS81/FIPS197
 - MD5/SHA-1/SHA256/384/512-96/128/192/256 with RFC1321 and FIPS180
 - HMAC in accordance with RFC2104/2403/2404 and FIPS198
 - ESN mode
 - GCM support for NSA "Suite B" (RFC6379/RFC6460) including GCM-128/256; GMAC-128/256
- Key Exchange Processor that supports high performance IKE and RSA computation
 - Public key exponentiation engine with hardware CRT support
 - Primary checking for RSA key generation
 - Handshake accelerator with automatic key material generation
 - True Random Number generator
 - Elliptic Curve support for NSA "Suite B"
 - Sub public key engine (PKCE) to support up to 4096 bit operation directly (4k for DH and 8k for RSA with CRT)
- DLP fingerprint support
 - TTTD (Two-Thresholds-Two-Divisors) content chunking
 - Two thresholds and two divisors are configurable

CP8 capabilities

The CP8 content processor provides the following services:

- Flow-based inspection (IPS, application control etc.) pattern matching acceleration
- High performance VPN bulk data engine
 - IPsec and SSL/TLS protocol processor
 - DES/3DES/AES in accordance with FIPS46-3/FIPS81/FIPS197
 - ARC4 in compliance with RC4
 - MD5/SHA-1/SHA256 with RFC1321 and FIPS180
 - HMAC in accordance with RFC2104/2403/2404 and FIPS198
- Key Exchange Processor support high performance IKE and RSA computation
 - Public key exponentiation engine with hardware CRT support
 - Primarily checking for RSA key generation
 - Handshake accelerator with automatic key material generation
 - Random Number generator compliance with ANSI X9.31
 - Sub public key engine (PKCE) to support up to 4096 bit operation directly
- Message authentication module offers high performance cryptographic engine for calculating SHA256/SHA1/MD5 of data up to 4G bytes (used by many applications)
- PCI express Gen 2 four lanes interface
- Cascade Interface for chip expansion

Kernel

Traffic is now in the process of exiting the FortiGate. The kernel uses the routing table to forward the packet out the correct exit interface.

The kernel also checks the NAT table and determines if the source IP address for outgoing traffic must be changed using SNAT. SNAT is typically applied to traffic from an internal network heading out to the internet. SNAT means the actual address of the internal network is hidden from the internet.

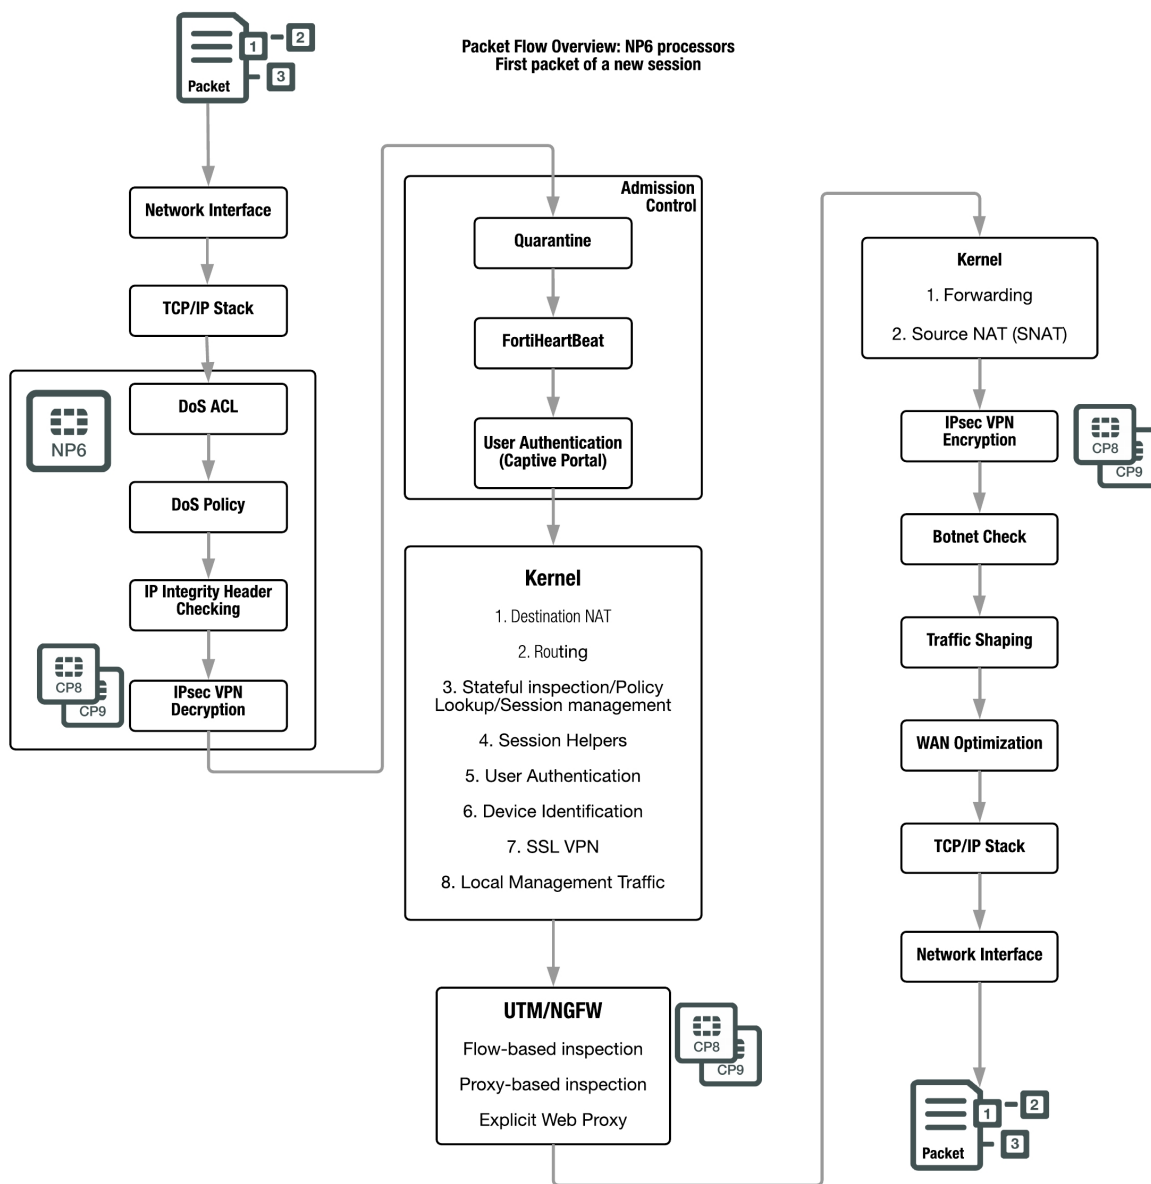
Egress

Before exiting the FortiGate, outgoing packets that are entering an IPsec VPN tunnel are encrypted and encapsulated. IPsec VPN encryption is offloaded to and accelerated by CP8 or CP9 processors. Packets are then subject to botnet checking to make sure they are not destined for known botnet addresses.

Traffic shaping is then imposed, if configured, followed by WAN Optimization. The packet is then processed by the TCP/IP stack and exits out the egress interface.

Packet flow: FortiGates with NP6 processors first packet of a new session

On a FortiGate with NP6 processors the first packet in a new session is handled the same way as on a FortiGate with no NP6 processors. Except that some processes, such as DoS, ACL, IP integrity checking, and IPsec VPN decryption are accelerated by the NP6 processor.



Network processors (NP6)

NP6 and NP6lite network processors provide fastpath acceleration by offloading communication sessions from the FortiGate CPU. When the first packet of a new session is received by an interface connected to an NP6 processor, just like any session connecting with any FortiGate interface, the session is forwarded to the FortiGate CPU where it is matched with a security policy. If the session is accepted by a security policy and if the session can be offloaded its session key is copied to the NP6 processor that received the packet. All of the rest of the packets in the session are intercepted by the NP6 processor and fast-pathed out of the FortiGate unit to their destination without ever passing through the FortiGate CPU. The result is enhanced network performance provided by the NP6 processor plus the network processing load is removed from the CPU. In addition the NP6 processor can handle some CPU intensive tasks, like IPsec VPN encryption/decryption.



NP6lite processors have the same architecture and function in the same way as NP6 processors. All of the descriptions of NP6 processors in this document can be applied to NP6lite processors except where noted.

Session keys (and IPsec SA keys) are stored in the memory of the NP6 processor that is connected to the interface that received the packet that started the session. All sessions are fast-pathed and accelerated, even if they exit the FortiGate unit through an interface connected to another NP6. There is no dependence on getting the right pair of interfaces since the offloading is done by the receiving NP6.

The key to making this possible is an Integrated Switch Fabric (ISF) that connects the NP6s and the FortiGate unit interfaces together. Many FortiGate units with NP6 processors also have an ISF. The ISF allows any port connectivity. All ports and NP6s can communicate with each other over the ISF. There are no special ingress and egress fast path requirements as long as traffic enters and exits on interfaces connected to the same ISF.

Some FortiGate units, such as the FortiGate-1000D include multiple NP6 processors that are not connected by an ISF. Because the ISF is not present fast path acceleration is supported only between interfaces connected to the same NP6 processor. Since the ISF introduces some latency, models with no ISF provide low-latency network acceleration between network interfaces connected to the same NP6 processor.


Each NP6 has a maximum throughput of 40 Gbps using 4 x 10 Gbps XAUI or Quad Serial Gigabit Media Independent Interface (QSGMII) interfaces or 3 x 10 Gbps and 16 x 1 Gbps XAUI or QSGMII interfaces.

There are at least two limitations to keep in mind:

- The capacity of each NP6 processor. An individual NP6 processor can support between 10 and 16 million sessions. This number is limited by the amount of memory the processor has. Once an NP6 processor hits its session limit, sessions that are over the limit are sent to the CPU. You can avoid this problem by as much as possible distributing incoming sessions evenly among the NP6 processors. To be able to do this you need to be aware of which interfaces connect to which NP6 processors and distribute incoming traffic accordingly.
- The NP6 processors in some FortiGate units employ NP direct technology that removes the ISF. The result is very low latency but no inter-processor connectivity requiring you to make sure that traffic to be offloaded enters and exits the FortiGate through interfaces connected to the same NP processor.

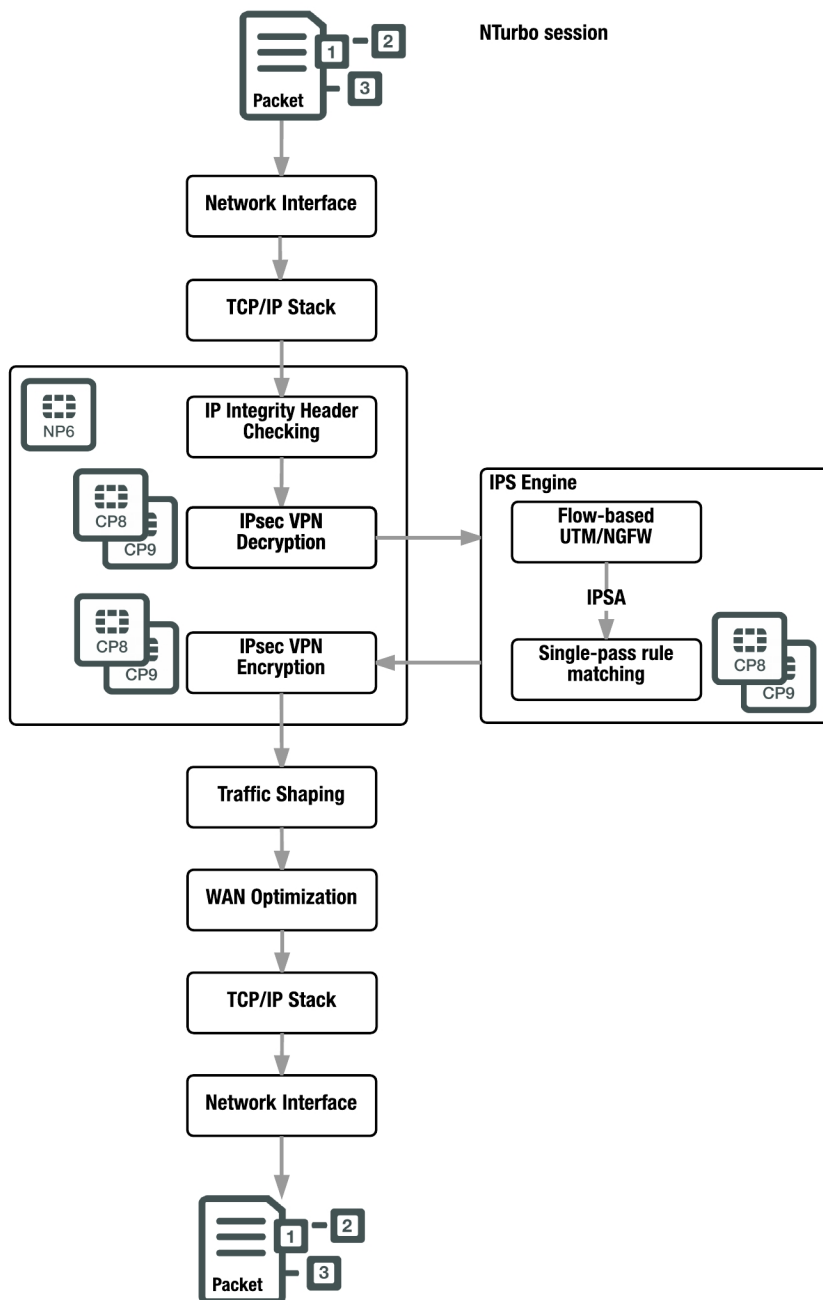
Packet flow: FortiGates with NP6 processors - packets in an offloaded session

The first packet of a session determines if the session can be offloaded. As long as there is no proxy-based UTM/NGFW, if your FortiGate includes NP6 processors, most sessions can be offloaded to them. After the first packet, subsequent packets in an offloaded session skip routing, UTM/NGFW, and kernel processors and are just forwarded out the egress interface by the NP6 processor. As well, security measures such as DoS policies, ACL, and so on are accelerated by the NP6 processor.



Packet flow: FortiGates with NP6 processors - packets in an NTurbo session

If your FortiGate supports NTurbo, many flow-based UTM/NGFW sessions can be offloaded to NP6 processors.



After the first packet, subsequent packets in an offloaded flow-based UTM/NGFW session skip routing, and kernel processors. Flow-based UTM/NGFW operations are still handled by the CPU with IPSA offloading pattern matching to CP8 or CP9 processors.

If a security threat is found the session is dropped. Otherwise, packets that are not blocked by UTM/NGFW are forwarded out of the egress interfaces by the NP6 processor.

NTurbo is not compatible with DoS policies, session helpers, or and most types of tunneling. If any of these features are present, flow-based UTM/NGFW sessions are not offloaded by NTurbo.

UTM/NGFW packet flow: flow-based inspection

Flow-based UTM/NGFW inspection identifies and blocks security threats in real time as they are identified using single-pass architecture that involves Direct Filter Approach (DFA) pattern matching to identify possible attacks or threats.

If a FortiGate or a VDOM is configured for flow-based inspection, depending on the options selected in the firewall policy that accepted the session, flow-based inspection can apply **IPS**, **Application Control**, **Web Filtering**, **DLP**, and **AntiVirus**. Flow-based inspection is all done by the IPS engine and as you would expect, no proxying is involved.



Flow-based DLP is supported but not recommended. Flow-based DLP is not available from the GUI, but can be configured from the CLI.

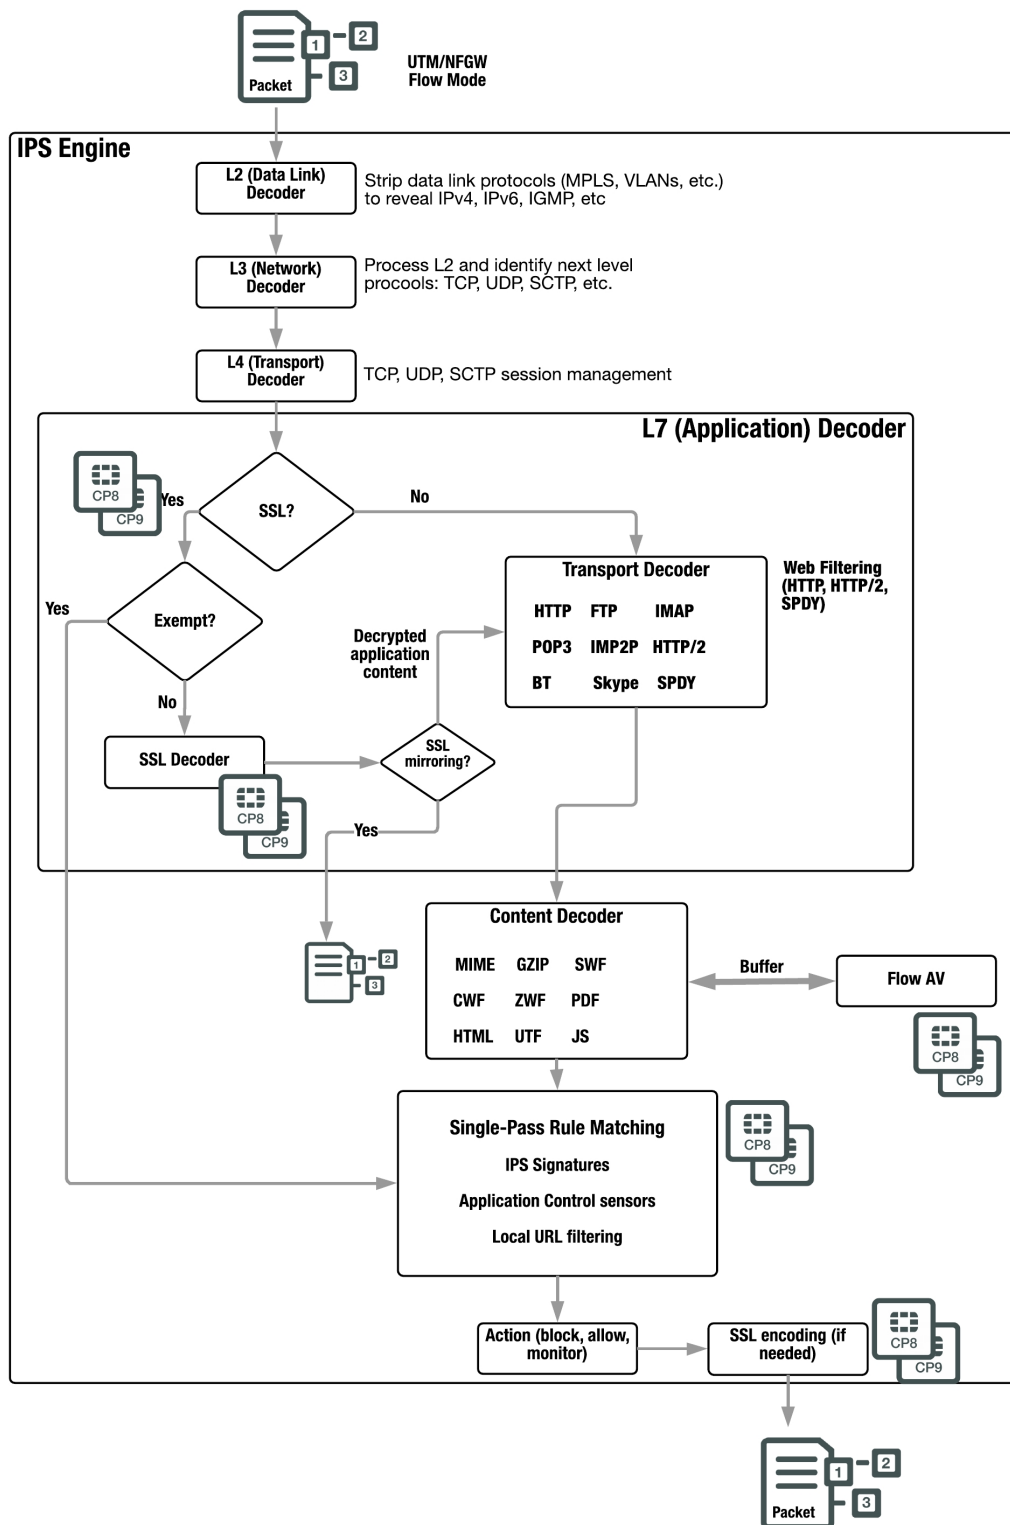
Before flow-based inspection can be applied the IPS engine uses a series of decoders to determine the appropriate security modules to be applied depending on the protocol of the packet and on policy settings. In addition, if SSL inspection is configured, the IPS engine also decrypts SSL packets. SSL decryption is offloaded and accelerated by CP8 or CP9 processors.

If your configuration includes SSL mirroring, the IPS engine copies decrypted application content, wraps it inside a TCP packet (with IP and ethernet headers), and sends it to the configured mirror interface. The TCP connection tuple is carried over from the original SSL connection. For the Ethernet frame, destination address is broadcast (FF:FF:FF:FF:FF:FF) and source address is all zeros.

All of the applicable flow-based security modules are applied simultaneously in one single pass, and pattern matching is offloaded and accelerated by CP8 or CP9 processors. IPS, Application Control, flow-based Web Filtering and flow-based DLP filtering happen together. Flow-based antivirus caches files during protocol decoding and submits cached files for virus scanning while the other matching is carried out.

Flow-based inspection typically requires less processing resources than proxy-based inspection and since its not a proxy, flow-based inspection does not change packets (unless a threat is found and packets are blocked). Flow-based inspection cannot apply as many features as proxy inspection (for example, flow-based inspection does not support client comforting and some aspects of replacement messages).

IPS and Application Control are only applied using flow-based inspection. Web Filtering, DLP and Antivirus can also be applied using proxy-based inspection.



UTM/NGFW packet flow: proxy-based inspection

If a FortiGate or VDOM is configured for proxy-based inspection then a mixture of flow-based and proxy-based inspection occurs. Packets initially encounter the IPS engine, which uses the same steps described in [UTM/NGFW packet flow: flow-based inspection on page 2298](#) to apply single-pass IPS and Application Control if configured in the firewall policy accepting the traffic.

The packets are then sent to the FortiOS UTM/NGFW proxy for proxy-based inspection. The proxy first determines if the traffic is SSL traffic that should be decrypted for SSL inspection. SSL traffic to be inspected is decrypted by the proxy. SSL decryption is offloaded to and accelerated by CP8 or CP9 processors.

If your configuration includes SSL mirroring, the IPS engine copies decrypted application content, wraps it inside a TCP packet (with IP and ethernet headers), and sends it to the configured mirror interface. The TCP connection tuple is carried over from the original SSL connection. For the Ethernet frame, destination address is broadcast (FF:FF:FF:FF:FF:FF) and source address is all zeros.

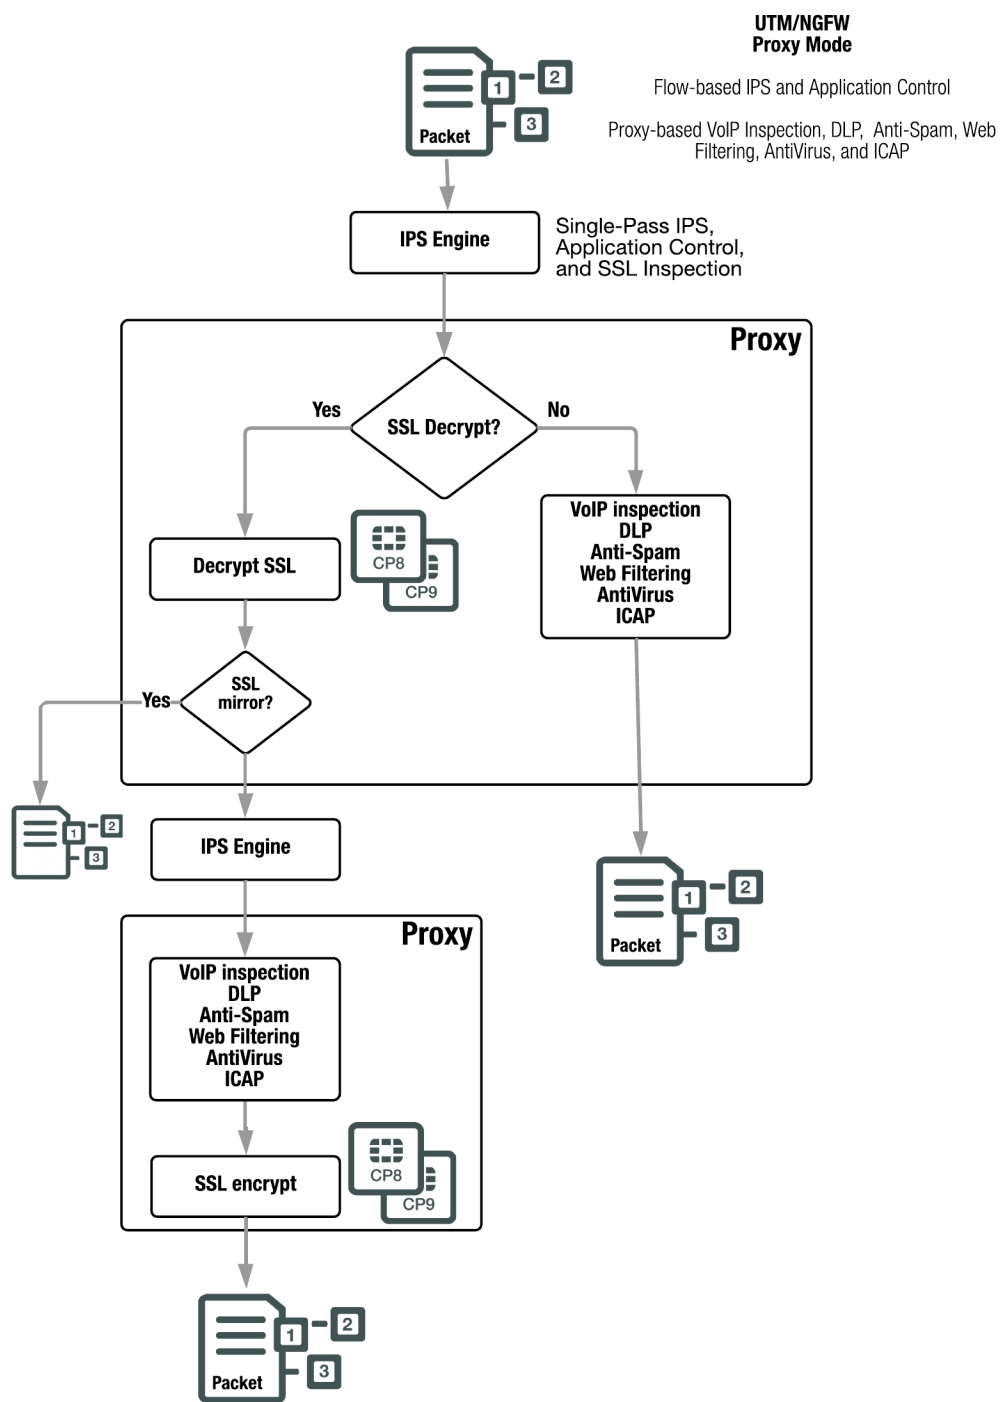
Proxy-based inspection extracts and caches content, such as files and web pages, from content sessions and inspects the cached content for threats. Content inspection happens in the following order: **VoIP inspection**, **DLP**, **Ant-Spam**, **Web Filtering**, **AntiVirus**, and **ICAP**.

If no threat is found the proxy relays the content to its destination. If a threat is found the proxy can block the threat and replace it with a replacement message.

Decrypted SSL traffic is sent to the IPS engine (where IPS and Application Control can be applied) before re-entering the proxy where actual proxy-based inspection is applied to the decrypted SSL traffic. Once decrypted SSL traffic has been inspected it is re-encrypted and forwarded to its destination. SSL encryption is offloaded to and accelerated by CP8 or CP9 processors. If a threat is found the proxy can block the threat and replace it with a replacement message.

The proxy can also block VoIP traffic that contains threats. VoIP inspection can also look inside VoIP packets and extract port and address information and open pinholes in the firewall to allow VoIP traffic through.

ICAP intercepts HTTP and HTTPS traffic and forwards it to an ICAP server. The FortiGate is the surrogate, or “middle-man”, and carries the ICAP responses from the ICAP server to the ICAP client; the ICAP client then responds back, and the FortiGate determines the action that should be taken with these ICAP responses and requests.



UTM/NGFW packet flow: explicit web proxy

If the explicit web proxy is enabled on a FortiGate or VDOM, a mixture of flow-based and proxy-based inspection occurs. One or more interfaces configured to listen for web browser sessions on the configured explicit web proxy port (by default 8080) accept all HTTP and HTTPS sessions on the explicit proxy port that match an explicit web proxy policy.

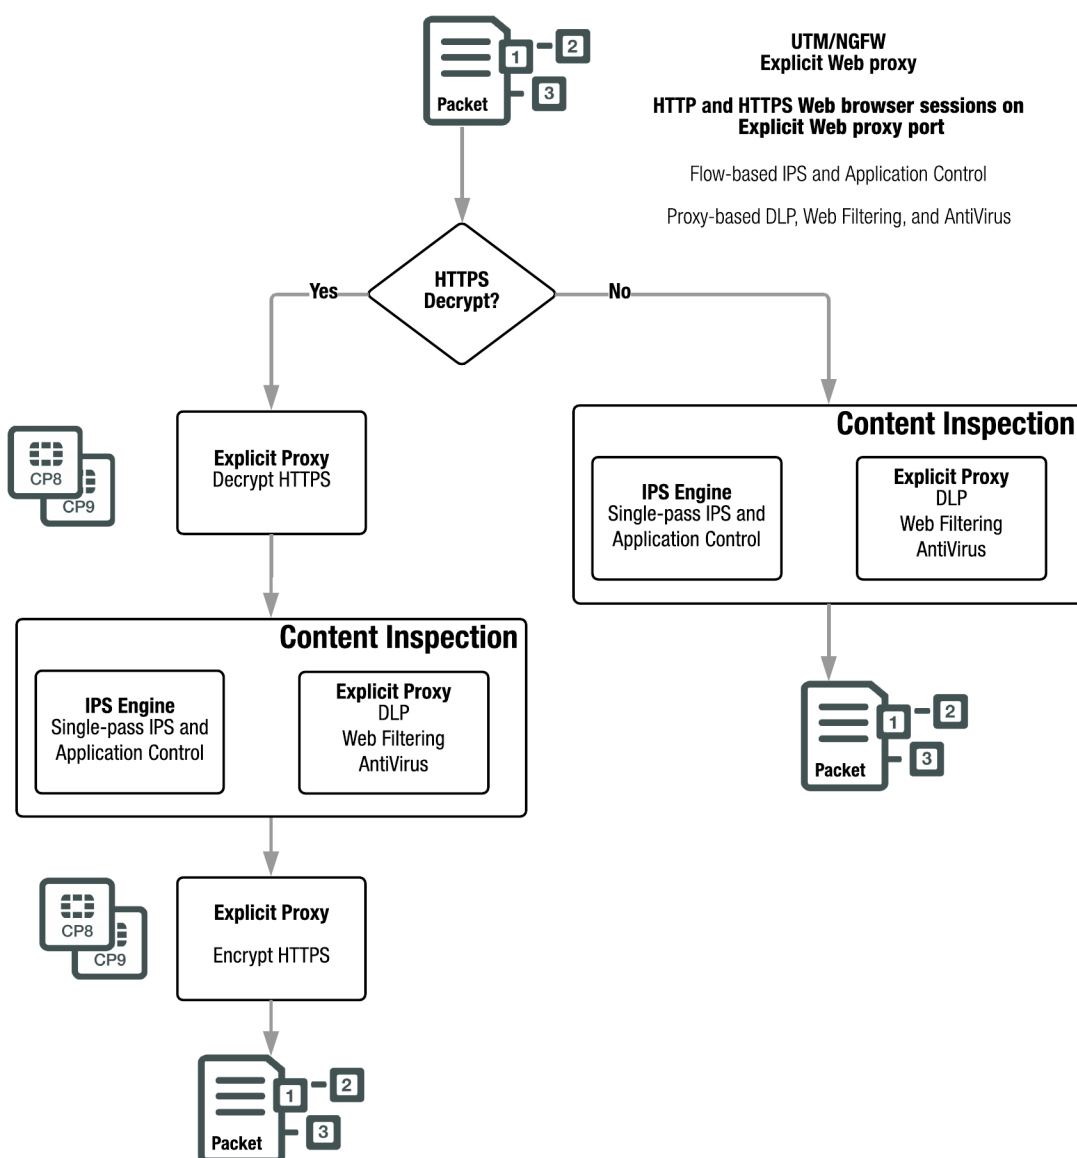
Plain text explicit web proxy HTTP traffic passes in parallel to both the IPS engine and the explicit web proxy for content scanning. The IPS engine applies IPS and application control content scanning. The explicit web proxy applies DLP, web filtering, and AntiVirus content scanning.

If the IPS engine and the explicit proxy do not detect any security threats, FortiOS relays the content to a destination interface. If the IPS engine or the explicit proxy detect a threat, FortiOS can block the threat and replace it with a replacement message.

Encrypted explicit web proxy HTTPS traffic passes to the explicit web proxy for decryption. Decrypted traffic once again passes in parallel to the IPS engine and the explicit web proxy for content scanning.

If the IPS engine and the explicit proxy do not detect any security threats, the explicit proxy re-encrypts the traffic and FortiOS relays the content to its destination. If the IPS engine or the explicit proxy detect a threat, FortiOS can block the threat and replace it with a replacement message. The explicit proxy offloads HTTPS decryption and encryption to CP8 or CP9 processors.

FortiOS uses routing to route explicit web proxy sessions through the FortiGate to a destination interface. Before a session leaves the exiting interface, the explicit web proxy changes the source addresses of the session packets to the IP address of the exiting interface. A FortiGate operating in transparent mode changes the source address to the transparent mode management IP address. You can also configure the explicit web proxy to keep the original client IP address.



Comparison of inspection types

The tables in this section show how different security functions map to different inspection types.

Mapping security functions to inspection types

The table below lists FortiOS security functions and shows whether they are applied by the kernel, flow-based inspection or proxy-based inspection.

FortiOS security functions and inspection types

Security Function	Kernel (Stateful inspection)	Flow-based inspection	Proxy-based inspection
Firewall	yes		
IPsec VPN	yes		
Traffic Shaping	yes		
User Authentication	yes		
Management Traffic	yes		
SSL VPN	yes		
IPS		yes	
Antivirus		yes	yes
Application Control		yes	
Web filtering		yes	yes
DLP		yes	yes
Email Filtering			yes
VoIP inspection			yes
ICAP			yes

More information about inspection methods

The three inspection methods each have their own strengths and weaknesses. The following table looks at all three methods side-by-side.

Inspection methods comparison

Feature	Stateful	Flow	Proxy
Inspection unit per session	first packet	selected packets, single pass architecture, simultaneous application of configured inspection methods	complete content, configured inspection methods applied in order
Memory, CPU required	low	medium	high
Level of threat protection	good	better	best
Authentication	yes		
IPsec and SSL VPN	yes		
Antivirus protection		yes	yes
Web Filtering		yes	yes
Data Leak Protection (DLP)		yes	yes
Application control		yes	
IPS		yes	
Delay in traffic	minor	no	small
Reconstruct entire content		no	yes

Chapter 19 - Sandbox Inspection

This guide explains how to set up sandbox inspection using FortiSandbox with a FortiGate. It contains the following sections:


- [An Overview of Sandbox Inspection](#): General information about how sandbox inspection works.
- [Using FortiSandbox with a FortiGate](#): How to set up sandbox inspection on a FortiGate.
- [Sandbox Integration](#): Integrating sandbox inspection with FortiGate, FortiSandbox, and FortiClient.
- [Sandbox Inspection FAQ](#): Frequently asked questions to help troubleshoot sandbox inspection.

An Overview of Sandbox Inspection

This section contains information about how Fortinet sandbox inspection works.

- [What is Sandbox Inspection?](#)
- [FortiSandbox Appliance vs FortiSandbox Cloud](#)
- [Sending Files for Sandbox Inspection](#)

What is Sandbox Inspection?



Sandbox inspection is a network process that allows files to be sent to a separate device, such as FortiSandbox, to be inspected without risking network security. This allows the detection of threats which may bypass other security measures, including zero-day threats.

You can configure your FortiGate device to send suspicious files to FortiSandbox for inspection and analysis. The FortiGate queries scan results and retrieves scan details. The FortiGate can also download malware packages as a complimentary AV signature database to block future appearances of the same malware and download URL packages as complimentary web filtering black list.

When a FortiGate sends files for sandbox inspection, the FortiSandbox uses virtual machines (VMs) running different operating systems to test the file and to determine if it is malicious. If the file exhibits risky behavior, or is found to contain a virus, a new signature can be added to the FortiGuard AntiVirus signature database.

When a FortiGate learns from FortiSandbox that a terminal is infected, the administrator can push instruction for self-quarantine on a registered FortiClient host.

FortiSandbox can process multiple files simultaneously since the FortiSandbox has a VM pool. The time to process a file depends on hardware and the number of sandbox VMs used to scan the file. It can take 60 seconds to five minutes to process a file.

FortiSandbox Appliance vs FortiSandbox Cloud

FortiSandbox is available as a physical or virtual appliance (FortiSandbox Appliance), or as a cloud advanced threat protection service integrated with FortiGate (FortiSandbox Cloud).

To select the settings for **Sandbox Inspection**, such as the FortiSandbox type, server, and notifier email, go to **Security Fabric > Settings**.

The table below highlights the supported features of both types of FortiSandbox:

Feature	FortiSandbox Appliance (including VM)	FortiSandbox Cloud
Sandbox inspection for FortiGate	Yes (FortiOS 5.0.4+)	Yes (FortiOS 5.2.3+)
Sandbox inspection for FortiMail	Yes (FortiMail OS 5.1+)	Yes (FortiMail OS 5.3+)
Sandbox inspection for FortiWeb	Yes (FortiWeb OS 5.4+)	Yes (FortiWeb OS 5.5.3+)
Sandbox inspection for FortiClient	Yes (FortiClient 5.4+ for Windows only)	No
Sandbox inspection for network share	Yes	No
Sandbox inspection for ICAP client	Yes	No
Manual File upload for analysis	Yes	Yes
Sniffer mode	Yes	Yes
File Status Feedback and Report	Yes	Yes
Dynamic Threat Database updates for FortiGate	Yes (FortiOS 5.4+)	Yes (FortiOS 5.4+)
Dynamic Threat Database updates for FortiClient	Yes (FortiClient 5.4 for Windows only)	Yes (FortiClient 5.6+ for Windows only)

Note that FortiMail keeps its own Dynamic Threat Database. For more information, see the [FortiSandbox documentation](#).

Sending Files for Sandbox Inspection

Sending files to the FortiSandbox appliance or to FortiSandbox Cloud does not block files immediately. Instead, the files assist in the discovery of new threats and the creation of new signatures to be added to the global FortiGuard AntiVirus database. Files deemed malicious are also immediately added to a custom Malware Package which is downloaded by the FortiGate every two minutes for live detection.

Enable **Sandbox Inspection** by going to **Security Fabric > Settings**. You can also configure the FortiSandbox type, server, and notifier email.

To see options for sending files for sandbox inspection, go to **Security Profiles > AntiVirus**. There are two options for sending files: **None** or **All Supported Files**. If **All Supported Files** is selected, users can withhold files from being submitted for inspection by type or name pattern.

To learn how to connect the FortiSandbox, go to ["Using FortiSandbox with a FortiGate" on page 2309](#)

Using FortiSandbox with a FortiGate

This section contains information about how to use sandbox inspection with FortiSandbox and FortiGate. It includes the following sections:

- [Connecting a FortiGate to FortiSandbox](#)
- [FortiSandbox Console](#)

Connecting a FortiGate to FortiSandbox

The procedures for connecting a FortiGate to FortiSandbox differ depending whether you are using [FortiSandbox Appliance](#) or [FortiSandbox Cloud](#).

If you are using FortiSandbox in a Fortinet Security Fabric, consult the [Fortinet Cookbook](#) site for the [Fortinet Security Fabric collection](#) of recipes.

Once the FortiGate is connected to FortiSandbox, an AntiVirus profile can be configured to send suspicious files for inspection. Sandbox integration can also be configured, for more information see "[Sandbox Integration](#)" on [page 2311](#).

Connecting to FortiSandbox Appliance

1. Connect the FortiSandbox Appliance to your FortiGate so that port 1 and port 3 on the FortiSandbox are on different subnets.



FortiSandbox port 3 is used for outgoing communication triggered by the execution of the files under analysis. While the FortiSandbox can accept files through any port, it is recommended to connect port 3 to a dedicated interface on your FortiGate to protect the rest of the network from threats currently being investigated by the FortiSandbox. Note too that port 1 can be used to accept files but is generally reserved for managing the FortiSandbox.

2. FortiSandbox port 3 must be able to connect to the Internet. On the FortiGate, go to **Policy & Objects > IPv4 Policy** and create a policy allowing connections from the FortiSandbox to the Internet (using the isolated interface on the FortiGate mentioned above). On FortiSandbox, network settings for port3 can be configured by going to **Scan Policy > General**.
3. On the FortiSandbox, go to **Network > System Routing** and add static routes for port 1.
4. On the FortiSandbox, go to **Dashboard** and locate the **System Information** widget. Now that the FortiSandbox has Internet access, it can activate its VM licenses. Wait until a green arrow shows up beside **Windows VM** before continuing to the next step.
5. On the FortiGate, go to **Security Fabric > Settings**. Select **Enable Sandbox Inspection** and select **FortiSandbox Appliance**. Set the **IP Address** and enter a **Notifier Email**. If you select **Test Connectivity**, the **Status** shows as **Service is not configured** because the FortiGate has not been authorized to connect to the FortiSandbox.
6. On the FortiSandbox, go to **Scan Input > Device**. **Edit** the entry for the FortiGate. Under **Permissions & Policy > Authorized**, select the checkbox and click **OK** to authorize the FortiGate.

- On the FortiGate, go to **Security Fabric > Settings** and select **Test Connectivity** for the FortiSandbox. The **Status** now shows that **Service is online**.

Connecting to FortiSandbox Cloud

Before you can connect a FortiGate to FortiSandbox Cloud, you need an active FortiCloud account. For more information, see the [FortiCloud documentation](#).

Once you have created a FortiCloud account, sandbox inspection should be enabled by default. To verify this, go to **Security Fabric > Settings**, enable **Sandbox Inspection**, and set to **FortiSandbox Cloud**.

To see the results from FortiSandbox Cloud in the FortiGate logs, go to **Log & Report > Log Settings** and enable **Send Logs to FortiCloud** and set **GUI Preferences** is to display logs from FortiCloud.

FortiSandbox Console

The FortiSandbox console is available at **FortiView > FortiSandbox**. The console displays all samples submitted for inspection. Information on the console can be filtered by checksum, file name, result, source, status, and user name.

Add Filter		Files	Source	5 minutes	1 hour	24 hours
Source	File Name	Status	Submitted			
vickimartin (192.168.200.110)	Breakpoints.js	Clean	10/02/2015 09:40:00			
vickimartin (192.168.200.110)	Corp_Reverb.css	Clean	10/02/2015 09:40:00			
vickimartin (192.168.200.110)	FortiOS%205.2%20CLI_sx.js	Clean	10/02/2015 09:40:00			
vickimartin (192.168.200.110)	Language.js	Clean	10/02/2015 09:40:00			
vickimartin (192.168.200.110)	MadCapAll.js	Clean	10/02/2015 09:40:00			
vickimartin (192.168.200.110)	Slideshow.css	Clean	10/02/2015 09:40:00			
vickimartin (192.168.200.110)	Toc.js	Clean	10/02/2015 09:40:00			
vickimartin (192.168.200.110)	Toc_Chunk6.js	Clean	10/02/2015 09:40:00			
vickimartin (192.168.200.110)	Web.css	Clean	10/02/2015 09:40:00			

If you right-click on an entry, you can choose to **Drill Down to Details**, **Quarantine Source Address**, or **Quarantine FortiClient Device**.

Information about the FortiSandbox database and sandboxing statistics are available at **Security Fabric > Settings** once sandbox inspection is enabled. The **Advanced Threat Protection** dashboard widget shows you the number of files that your FortiGate unit has uploaded or submitted to FortiSandbox.

Refer to [FortiSandbox documentation](#) for details on what you can access through the FortiSandbox GUI .

Sandbox Integration

Sandbox integration adds another level to sandbox inspection, allowing you to set up automatic actions to protect your network from files FortiSandbox determines are malicious. These actions include: receiving AntiVirus signature updates from FortiSandbox, adding the originating URL of any malicious file to a blocked URL list, and extending sandbox scanning to FortiClient devices.

This section contains the following topics:

- [Overview](#)
- [Example Configuration](#)

Overview

FortiSandbox integration involves three different FortiGate security profiles: [AntiVirus](#), [Web Filtering](#), and [FortiClient Profiles](#).

A FortiGate can retrieve scan results and details from FortiSandbox, and also receive antivirus and web filtering signatures to supplement the current signature database. When FortiGate learns from FortiSandbox that an endpoint is infected, the administrator can push instruction for self-quarantine on a registered FortiClient host.

When integrated with a FortiGate unit, the following protocols are supported by FortiSandbox: HTTP, HTTPS, FTP, FTPS, POP3, POP3S, IMAP, IMAPS, SMTPS, MAPI, MAPIS, SMB, and supported IM protocols.

AntiVirus

When FortiSandbox discovers a malicious file, it can create an AntiVirus signature for that file and add that signature to both the local FortiGate malware database and the FortiGuard AntiVirus signature database. Through FortiSandbox integration, this signature can be sent to a FortiGate to block the file from re-entering the network and to prevent the future retransmission of that file to FortiSandbox.

Use of the FortiSandbox AntiVirus database is enabled in an AntiVirus profile, found at **Security Profiles > AntiVirus**. It can also be configured using the following CLI commands:

```
config antivirus profile
  edit <profile>
    set analytics-db enable
  end
```

Web Filtering

FortiSandbox integration can also be used to allow FortiSandbox to add a URL filter blocking the source of a discovered malicious file to the FortiGate's blocked URL list.

Blocking malicious URLs discovered by FortiSandbox is enabled in a Web Filter profile, found at **Security Profiles > Web Filter**. It can also be configured using the following CLI commands:

```
config webfilter profile
  edit <profile>
    config web
      set blacklist enable
    end
```

FortiClient Profiles



Extended FortiSandbox scanning is currently only supported by FortiClient 5.4 for Windows. It can also only be used with FortiSandbox Appliance.

When extended FortiSandbox scanning is enabled for FortiClient, files downloaded by FortiClient can be sent to the FortiSandbox for inspection. Also, if a suspicious file is discovered, FortiClient can be configured to wait until sandbox inspection is complete before allowing that file to be accessed.

AntiVirus signatures can also be pushed by the FortiGate to FortiClient.

If a FortiClient device attempts to download a file that FortiSandbox discovers is malicious, the FortiSandbox notifies the FortiGate. The administrator can take action to quarantine the device. When a quarantine is in effect, FortiClient cuts off other network traffic from the device directly, preventing it from infecting or scanning the local network. When a device is under quarantine, FortiClient cannot be shutdown or uninstalled. A user is also unable to unregister from the FortiGate that quarantined them, or register to another FortiGate unit. A quarantine can only be lifted by the administrator of the FortiGate where the FortiClient device is registered.

Extending FortiSandbox scanning can be configured in the **Security** settings of a FortiClient Profile, found at **Security Profiles > FortiClient Compliance**. It can also be configured using the following CLI commands:

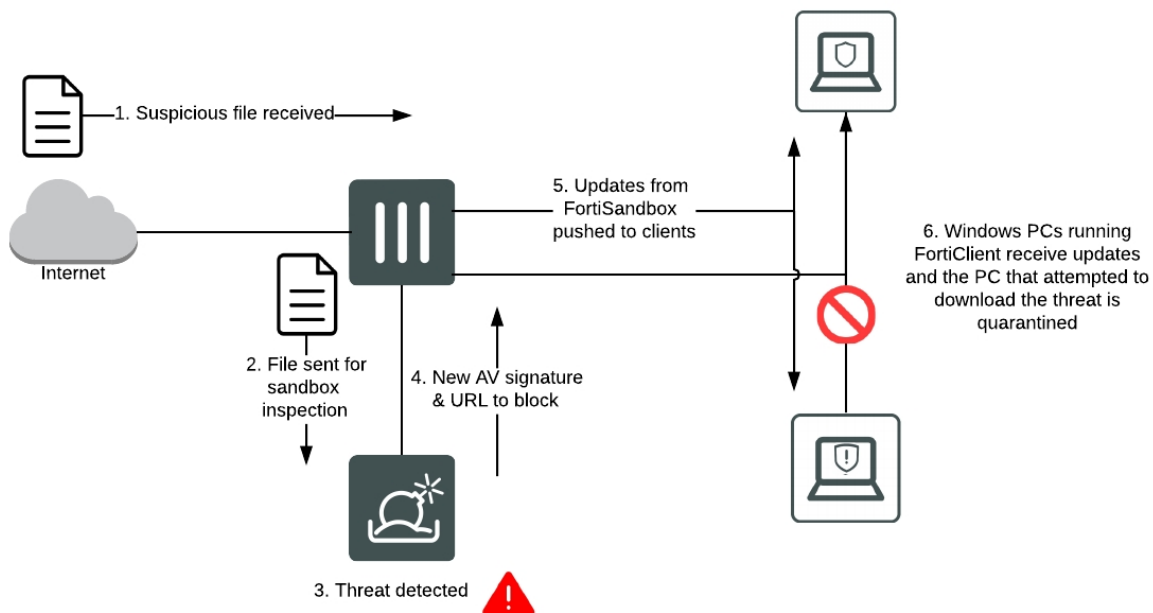
```
config endpoint-control profile
edit <profile>
config forticlient-winmac-settings
set forticlient-av enable
set av-realtime-protection enable
set sandbox-analysis enable
set sandbox-address <address>
end
```

Extending FortiSandbox scanning can also be configured directly in the FortiClient **AntiVirus** settings. If you are using FortiClient version 5.6+, the **Sandbox Detection** feature can be used to send files to FortiSandbox for analysis without having to install the AntiVirus feature. See the FortiClient 5.6 Administration Guide for details.

The number of files sent from a single device to FortiSandbox can be limited by [configuring the submission limit](#) on the FortiSandbox. This allows users to prioritize which devices get the greater share of FortiSandbox resources.

Example Configuration

The following example configuration sets up FortiSandbox integration using AntiVirus, Web Filtering, and a FortiClient profile. This configuration assumes that a connection has already been established between the FortiSandbox Appliance and the FortiGate.



1. Go to **Security Fabric > Settings** and confirm that **Sandbox Inspection** is enabled and the **FortiSandbox Appliance** is connected.
2. Go to **Security Profiles > AntiVirus** and edit the default profile. Under **Inspection Options**, select **All Supported Files** to be sent for inspection and enable **Use FortiSandbox Database**. You have the option of withholding files by name or pattern. Select **Apply**.
3. Go to **Security Profiles > Web Filter** and edit the default profile. Under **Static URL Filter**, enable **Block malicious URLs discovered by FortiSandbox**. Select **Apply**.
4. Go to **Security Profiles > FortiClient Compliance** and edit the default profile. Under **Security Posture Check**, enable **Realtime Protection**. Next, enable **Scan with FortiSandbox**. Select **Apply**.
5. Go to **Policy & Objects > IPv4 Policy** and view the policy list. If a policy has AntiVirus and Web Filtering profiles scanning applied, the profiles will be listed in the **Security Profiles** column. If scanning needs to be added to any security policy (excluding the **Implicit Deny** policy) select the **+** button in the **Security Profiles** column for that policy, then select the default **AntiVirus Profile**, the default **Web Filter Profile**, the appropriate **Proxy Options**, and select the **deep-inspection** profile for **SSL/SSH Inspection** (to ensure that encrypted traffic is inspected).
6. Select **OK**.

Results

If your FortiGate discovers a suspicious file, it will be sent to the FortiSandbox. To view information about the files that have been sent on the FortiGate, go to **FortiView > FortiSandbox** to see a list of file names and current status.

To view results on the FortiSandbox, go to the **Dashboard** and view the **Scanning Statistics** widget. There may be a delay before results appear on the FortiSandbox.

Open FortiClient using a Windows PC on the internal network. Make sure it is registered to your FortiGate. Go to the **AntiVirus** tab and open **Settings**. You will see that the **Realtime Protection** settings match the FortiClient profile configured on the FortiGate. These settings cannot be changed using FortiClient.

If a PC running FortiClient downloads a suspicious file that the FortiSandbox determined was malicious, a quarantine would be applied automatically. While the quarantine is in effect, FortiClient cannot be shutdown on the PC. It can not be uninstalled or unregistered from the FortiGate. The quarantine can only be released from the FortiClient Monitor on the FortiGate.

Sandbox Inspection FAQ

The following are some frequently asked questions about using sandbox inspection with FortiSandbox and FortiGate.

Why is the FortiSandbox Cloud option not available when sandbox inspection is enabled?

This option is only available if you have already created a FortiCloud account. For more information, see the [FortiCloud documentation](#).

Why don't results from FortiSandbox Cloud appear in the FortiGate GUI?

Go to **Log & Report > Log Settings** and make sure **Send Logs to FortiCloud** is enabled and **GUI Preferences** is set to **Display Logs from FortiCloud**.

Why are the FortiSandbox Appliance VMs inactive?

Make sure that port 3 on the FortiSandbox has an active Internet connection. This is required in order to active the FortiSandbox VMs.

Why aren't files are being scanned by FortiSandbox?

Make sure an AntiVirus profile that sends files to FortiSandbox is enabled for all policies that require sandbox inspection.

Is FortiSandbox supported by FortiGate when in NAT or Transparent mode?

Yes, both NAT and Transparent mode are supported.

Are FortiGates behind a NAT device supported? If so how many?

Yes, multiple FortiGates can be supported in-line with FortiSandbox. Currently, there is a limitation where the FortiSandbox will see all FortiGates only as one device so there is no way to differentiate reports but all material will be sent.

If the FortiGate has a dynamic IP, will the FortiSandbox automatically update the FortiGate?

Yes. Dynamic IPs™ are supported and the FortiGate will not have to be reconfigured on the FortiSandbox each time.

Chapter 20 - Fortinet Security Fabric

Introduction

The Fortinet Security Fabric is an end-to-end security solution that gives you control, integration, and easy management of security across your entire organization. The Security Fabric provides an intelligent architecture that interconnects discrete security solutions into an integrated whole to detect, monitor, block, and remediate attacks across the entire enterprise attack surface.

This document is a complete reference guide for the Security Fabric, including an overview of what the Security Fabric is, what devices are included in the Security Fabric and how they work together to secure your network, and how to configure and manage the Security Fabric.

What's new in FortiOS 6.0.2

The following list contains new Fortinet Security Fabric features added in FortiOS 6.0.2:

Automation features

- You can now test automation stitches in the GUI or by using the `diagnose automation test` command.
- When an automation stitch is triggered, the FortiGate creates an event log.

For more information, see ["Automation stitches" on page 2349](#).

Fabric Connector features

- FortiOS supports SDN Connectors to Google Cloud Platform (GCP).
- SDN Connectors to Amazon Web Services no longer require setting a VPC ID.
- SDN Connectors to Microsoft Azure no longer require setting an Azure subscription ID or Azure resource group.
- You can set the region for Azure to the Germany Azure Server or the US Government Azure Server.

For more information, see ["Fabric Connectors" on page 2362](#).

What's new in FortiOS 6.0.1

The following list contains new Fortinet Security Fabric features added in FortiOS 6.0.1:

New features

- Chaining and delaying actions for AWS Lambda and webhook
- Diagnose commands for automation stitches
- Available services for Fabric Connectors
- Verifying Fabric Connector status
- FortiManager in backup mode

6.0.1 GUI changes

- The options to configure single sign-on are now located at **Security Fabric > Fabric Connectors**.
- Automation stitches are available for FortiGate devices that don't belong to a Security Fabric.
- Two new triggers are available: **Security Rating Summary** and **AV & IPS DB Update**.
- The **Event Log** trigger now shows a list of events that can be used as triggers.
- The **FortiExplorer Notification** action now has warning and information message options.
- The **IOC level threshold** option is now called the **Threat level threshold**.
- The triggers for **Conserve Mode** and **High CPU** are now available only in the CLI.

What's new in FortiOS 6.0.0

The following list contains new Fortinet Security Fabric features added in FortiOS 6.0:

- [Automation stitches](#)
- [Security Fabric Rating license](#)
- [To authorize serial numbers of devices from the root FortiGate](#)
- [Joining the Security Fabric by device request](#)
- For Physical and Logical Topology enhancements, see:
 - [The WAN Cloud icon](#)
 - [Switch stacking](#)
 - [FortiAP and FortiSwitch integrations](#)
 - [Distinguishing client traffic from server traffic](#)
 - [Using the Search bar to find information in the topology views](#)
- [FortiMail Stats widget](#)
- [Desynchronizing the FortiAnalyzer, FortiSandbox, and FortiManager](#)
- [Fabric Connectors](#)

Fortinet Security Fabric overview

The Fortinet Security Fabric provides a visionary approach to security that allows your organization to deliver intelligent, powerful, and seamless security. Fortinet offers security solutions for endpoints, access points, network elements, the data center, applications, cloud, and data, designed to work together as an integrated Security Fabric that can be integrated, analyzed, and managed to provide end-to-end protection for your network. Your organization can also add third-party products that are members of the Fabric-Ready Partner Program to the Security Fabric.



All elements in the Security Fabric work together as a team to share policy, threat intelligence, and application flow information. This collaborative approach expands network visibility and provides fast threat detection in real time and the ability to initiate and synchronize a coordinated response, no matter which part of the network is being compromised. The Security Fabric allows your network to automatically see and dynamically isolate affected devices, partition network segments, update rules, push out new policies, and remove malware.

The Security Fabric is designed to cover the entire attack surface and provide you with complete visibility into your network. It allows you to collect, share, and correlate threat intelligence between security and network devices,

centrally manage and orchestrate policies, automatically synchronize resources to enforce policies, and coordinate a response to threats detected anywhere across the extended network. The unified management interface provides you with cooperative security alerts, recommendations, audit reports, and full policy control across the Security Fabric that will give you confidence that your network is secure.

Access security

The Security Fabric secures the access layer of your organization's network. It integrates various access points in a network, such as endpoints, applications, the cloud, and IoT devices, regardless of their distribution, into an end-to-end solution that covers all attack surfaces.

Secure access architecture extends coordinated security policies to the edge of the wired and wireless network, where most vulnerabilities are targeted. It protects the access layer, guarding against data breaches and security threats from both internal user devices and IoT products.

Client security

Client security, through FortiClient, provides easy-to-manage, automated, fully customizable endpoint security for various devices. FortiClient provides end-to-end threat visibility and control by natively integrating endpoints into the security architecture and offers unified endpoint features, including compliance, protection, and secure access. It also offers integrated patch management and vulnerability shielding to harden all endpoints.

FortiClient integrates with the Security Fabric to provide real-time actionable visibility to stop threats to your organization's network at the endpoints.

For more information about FortiClient, see <http://www.forticlient.com/>.

Application security

The Security Fabric protects your organization's sensitive and proprietary data that is managed by applications, and ensures the security and availability of your organization's applications. It allows Fortinet application security products, and those of third-party vendors, to work together to boost security across core networks, remote devices, and the cloud. This provides your organization with a network architecture that is secure, aware, actionable, scalable, and open.

Fortinet's robust and integrated application security solution provides a complete end-to-end high-performance solution that protects your organization's valuable information by using a combination of Fortinet products which are deeply integrated into the Security Fabric for direct communications. These products include web application firewalls for application security, DDoS attack mitigation appliances for DDoS protection, advanced application delivery controllers (ADCs) to meet the demands of secure application traffic, sandboxing to isolate malicious code for inspection, and email security gateways that can detect and prevent email-borne threats from getting to your users.

Cloud security

The Security Fabric is designed to extend deep into different cloud environments to ensure that policies are consistent and enforced across all distributed resources. Within the unified security architecture, virtual firewalls can be deployed across private, public, and hybrid clouds to establish north-south and east-west microsegmentation. The Security Fabric weaves cloud applications into the broader environment, governed by seamless, universal security and compliance policies and managed using transparent visibility across the entire attack surface. Combining Fortinet Cloud Security with an existing enterprise firewall deployment extends the

same powerful security, as well as the same intelligence and dynamic risk mitigation to applications located either in the cloud or on-premise.

NOC and SOC security

Fortinet's security operations center solution covers both IT and security risk management across your entire organization. The solution is a comprehensive approach to managing risk that includes adaptive awareness of the threat landscape, rapid local and global threat detection, reduced complexity in managing alerts and alarms, and reporting and analytics so you can better understand how your organization's risk profiles are being managed.

When Fortinet devices are unified into a Security Fabric, with compatible operating systems and shared intelligence, the security operations solution also includes information from network elements beyond Fortinet devices. The solution allows your network operations center (NOC) and security operations center (SOC) to share information, integrating and cross-correlating the data from each operations center. This additional context, visibility, and focus breaks down the barrier between your NOC and SOC, and gives you a comprehensive view across your entire Security Fabric so you can quickly find and respond to threats.

Advanced threat intelligence

Fortinet's Advanced Threat Protection (ATP) solution allows your organization to detect and mitigate against threats, both known and unknown, and share that information locally to deliver a coordinated defense.

The ATP solution relies on many types of security technologies, products, and research applied from the network edge through to endpoint devices. To deliver the most effective protection, they are integrated with other security elements from the Enterprise Firewall and Cloud solutions to work together automatically, continuously handing off data from one element to the next to identify, evaluate, and respond to attacks across the entire environment.

The ATP framework delivers end-to-end protection across the attack chain and consists of three elements: prevention, detection, and mitigation, with continuous threat monitoring and analytics from FortiGuard Labs.

For more information about the Advanced Threat Protection Solution, see <http://www.fortinet.com/atp>.

Partner API

The Fortinet Fabric-Ready Partner Program is an interoperability program for technology alliance partners. Technology alliance partners integrate their products with the Fortinet Security Fabric using Fortinet Security Fabric APIs. Their products are then able to actively collect and share threat and mitigation information from one end of the security solution to the other, which improves threat intelligence, enhances overall threat awareness, and broadens threat response.

Inclusion in the program means that Fabric-Ready Partners have collaborated with Fortinet and leveraged the Fortinet Security Fabric APIs to develop and validate integrated end-to-end security solutions that are ready for deployment.

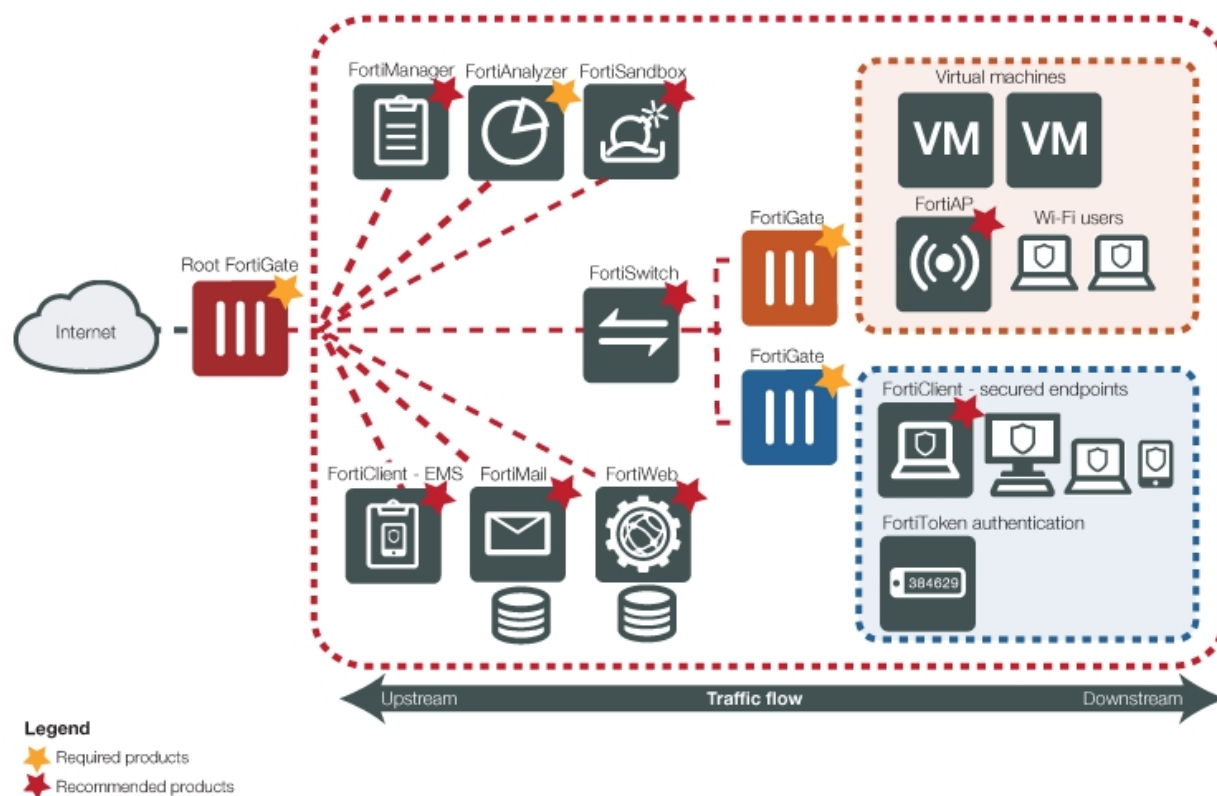
The Fabric-Ready Partner Program allows Fortinet technology alliance partners to build on Fortinet products and solutions which help your organization get even more value from your security deployment.

For more information about the Fortinet Fabric-Ready Partner Program, see

<https://www.fortinet.com/partners/partnerships/alliance-partners.html>

The Security Fabric solution components

The Fortinet Security Fabric consists of various components that work together to form the Security Fabric that secures your organization's network. The following diagram shows an example Security Fabric that contains both required and recommended Fortinet products:



Devices in the Security Fabric

The Security Fabric implementation consists of:

- Required devices
- Recommended devices
- Optional devices

Required devices

The following table shows devices that are required in the Fortinet Security Fabric:

Device	Description
FortiGate	<p>FortiGate is a next-generation firewall (NGFW) that provides enterprise-class protection against network, content, and application-level threats.</p> <p>FortiGate devices are the core of the Security Fabric and can have one of the following roles in the Security Fabric:</p> <ul style="list-style-type: none"> • Root FortiGate: The root FortiGate is the main component in the Security Fabric. It is typically located on the edge of the network and connects the internal devices and networks to the Internet through your ISP. From the root FortiGate, you can see information about the entire Security Fabric from the Physical and Logical Topology pages in the Security Fabric menu. • Internal Segmentation Firewall (ISFW): After a root FortiGate is installed, all other FortiGate devices in the Security Fabric act as ISFWs. An ISFW is a firewall that is located at strategic points in your internal network, rather than on the network edge. This allows extra security measures to be taken around key network components, such as servers that contain valuable intellectual property. ISFW FortiGate devices create network visibility by sending traffic and information about the devices that are connected to them to the root FortiGate.
FortiAnalyzer	<p>FortiAnalyzer collects, analyzes, and correlates log data from Fortinet devices throughout your organization's network, and allows you to view all firewall traffic and generate reports from a single console.</p> <p>FortiAnalyzer gives you increased visibility into your organization's network and simplifies network logging by storing and displaying all log information in one place. It provides centralized monitoring and awareness of threats, events, and network activity by collecting and correlating logs from Security Fabric devices, such as FortiGate, FortiClient, FortiSandbox, FortiWeb, and FortiMail. This gives you a deeper and more comprehensive view across your entire Security Fabric. You can use the robust security alert information and real-time threat intelligence that FortiAnalyzer provides to quickly identify and respond to security threats across your organization's network.</p>

Recommended devices

The following table shows devices that Fortinet recommends you have in the Fortinet Security Fabric:

Device	Description
FortiAP	<p>FortiAP is a wireless access point that provides integrated, secure, identity-driven wireless LAN access for your organization's network.</p> <p>You can add FortiAP devices to extend the Security Fabric to your wireless devices. Devices connected to a FortiAP appear in the Physical and Logical Topology pages in the Security Fabric menu.</p>

Device	Description
FortiClient	<p>FortiClient adds endpoint control to devices that are located in the Security Fabric, allowing only traffic from compliant devices to flow through the FortiGate. This is done through FortiClient compliance profiles.</p> <p>In the Security Fabric, FortiClient compliance profiles are applied by the first FortiGate that a device's traffic flows through. This is often an ISFW FortiGate. Device registration and on-net status information for a device that is running FortiClient appears only on the FortiGate that applies the FortiClient profile to the device.</p>
FortiClient EMS	<p>FortiClient Enterprise Management Server (EMS) is a security management solution that provides scalable and centralized management of multiple endpoint devices.</p> <p>FortiClient EMS is used in the Security Fabric to provide visibility across your network, to securely share information, and assign security profiles to endpoints.</p>
FortiMail	<p>FortiMail is a secure email gateway that uses various threat prevention methods, including antispam, antimalware, sandboxing, and anomaly detection.</p> <p>FortiMail integrates with other Fortinet products, as well as third-party virtual and cloud platforms, to help establish a seamless Security Fabric across the entire attack surface. FortiMail anti-spam processing helps offload other devices in the Security Fabric that would typically carry out this process.</p>
FortiManager	<p>FortiManager is an easy-to-use, single pane of glass management console, that gives you total visibility, full control, and complete protection of your organization's network.</p> <p>Using the FortiManager in the Security Fabric allows you to simplify the network management of devices in the Security Fabric by centralizing management access in a single device. This allows you to easily control the deployment of security policies, FortiGuard content security updates, firmware revisions, and individual configurations for devices in the Security Fabric.</p>
FortiSandbox	<p>FortiSandbox is an advanced threat protection appliance that improves your security architecture by identifying and validating threats in a separate, secure environment.</p> <p>You can add FortiSandbox to your Security Fabric to improve security with sandbox inspection. Sandbox integration allows FortiGate devices in the Security Fabric to automatically receive signature updates from FortiSandbox and add the originating URL of any malicious file to a blocked URL list.</p>

Device	Description
FortiSwitch	<p>FortiSwitch is a secure access switch that can be integrated into the Fortinet Security Fabric through the FortiLink protocol. FortiLink allows FortiSwitch ports to become logical extensions of the FortiGate. This allows the FortiGate to auto-discover a connected FortiSwitch for provisioning, including the attachment of policy to ports or VLANs. With an integrated access layer, the FortiGate provides consolidated visibility and reporting with Physical and Logical Topology views of the Security Fabric in the Security Fabric menu.</p> <p>You can add a FortiSwitch to the Security Fabric when it is managed by a FortiGate within the Security Fabric, and connected to an interface that uses FortiTelemetry.</p> <p>Devices connected to the FortiSwitch appear in the Physical and Logical Topology pages in the Security Fabric menu, and security features, such as FortiClient compliance profiles, are applied to them.</p>
FortiWeb	<p>FortiWeb is a web application firewall that protects hosted web applications from attacks that target known and unknown exploits.</p> <p>In the Security Fabric, FortiWeb defends the application attack surface from attacks that target application exploits. You can also configure FortiWeb to apply web application firewall features, virus scanning, and web filtering to HTTP traffic to help offload other devices in the Security Fabric that would typically carry out these processes.</p>

Optional devices

The following table shows devices that are optional in the Fortinet Security Fabric:

Device	Description
Other Fortinet products	Many other Fortinet products can be added to the Security Fabric, including FortiAuthenticator, FortiToken, FortiCache, and FortiSIEM.
Third-party products	Third-party products that belong to the Fortinet Fabric-Ready Partner Program .

Security Fabric topology views

You can see the Security Fabric topology in the root FortiGate GUI. Two viewing options are available: the Physical Topology view and the Logical Topology view.

The Physical Topology view displays the physical structure of your network, by showing the devices in the Security Fabric and the connections between them. The Logical Topology view displays the logical structure of your network, by connection, by showing information about logical and physical network interfaces in the Security Fabric and the interfaces that connect devices in the Security Fabric. Only Fortinet devices are shown in the topology views.

For more information about the topology views, see ["Using the Fortinet Security Fabric" on page 2337](#).

Security Fabric Rating

The Security Fabric Rating provides a method to continually monitor and improve your organization's Security Fabric configuration. The Security Fabric Rating is a feature on the FortiGate that analyzes your Security Fabric deployment, identifies potential vulnerabilities, and highlights best practices that you can use to improve the overall security and performance of your network.

Using the Security Fabric Rating helps you to:

- Tune your network configuration
- Deploy new hardware and software
- Have more visibility into your network
- Gain more control over your network
- Adhere to your organization's compliance requirements

The Security Fabric Rating provides a Security Fabric Score based on how many security checks your network passes and fails during the test. By checking the Security Fabric Score, and implementing the recommendations, you can have confidence that your network is getting more secure over time.

For more information about running a Security Fabric Rating check, see ["Using the Fortinet Security Fabric" on page 2337](#).

FortiTelemetry

FortiTelemetry is a protocol that Fortinet products in the Security Fabric use to communicate with each other. It connects Security Fabric devices and allows dynamic status updates to travel between them. The Security Fabric uses FortiTelemetry to link various security sensors and tools together to collect, coordinate, and respond to malicious behavior anywhere it occurs in your network in real time.

You must enable FortiTelemetry on interfaces that connect Fortinet devices in the Security Fabric.

Configuring the Fortinet Security Fabric

This section contains information about how to configure a Fortinet Security Fabric:

- [Forming the Security Fabric](#)
- [Setting up data collection with FortiAnalyzer](#)
- [Adding a FortiSandbox to the Security Fabric](#)
- [Adding a FortiManager to the Security Fabric](#)
- [Adding FortiClient EMS to the Security Fabric](#)

System requirements

To set up the Security Fabric in FortiOS 6.0, the devices that you want to include in the Security Fabric must meet the Product Integration and Support requirements in the [FortiOS Release Notes](#).

Some features of the Security Fabric are available only in certain firmware versions and models. Not all FortiGate models can run the FortiGuard Security Rating Service if they are the root FortiGate in a Security Fabric. For more information, see the Special Notices in the [FortiOS Release Notes](#).

For more information about upgrading the Security Fabric to version 6.0, see the [Fortinet Security Fabric Upgrade guide](#).

Prerequisites

- Determine which devices you want to have in the Security Fabric.
- Ensure devices meet the [Configuring the Fortinet Security Fabric](#)
- If devices are not already installed in your network, complete basic installation and configuration tasks by following the instructions in the device documentation.
- Disable virtual domains (VDOMs) on all FortiGate devices that you want to add to the Security Fabric.
- Configure all FortiGate devices that you want to add to the Security Fabric to use NAT/route mode.

Forming the Security Fabric

To form the Security Fabric, you configure the root FortiGate and then the ISFW FortiGate devices. Although you can configure any of the FortiGate devices in the Security Fabric to be the root FortiGate, you typically configure the edge FortiGate as the root FortiGate. This setup allows you to view the full topology of the Security Fabric from the top down.

The following procedures include configuration steps for a typical Security Fabric implementation, where the root FortiGate is the edge FortiGate and the ISFW FortiGate devices are all FortiGate devices that are downstream from the root FortiGate.

Adding devices to the Security Fabric

You can easily and securely allow FortiGate, FortiAP and FortiSwitch to join the Security Fabric without sharing the password of the root FortiGate. You can authorize these device serial numbers from the root FortiGate or allow the device to join by request. New authorization requests include the serial number of the device, the device IP address, and a list of High Availability (HA) members.

HA members can contain up to four serial numbers and this list is used to ensure that, in the event of failover, the secondary FortiGate is still authorized.

After a FortiGate or FortiWiFi joins the Security Fabric, any connected FortiAP or FortiSwitch automatically appears in the topology. You can then authorize these additional devices from the FortiGate or FortiWiFi they're connected to or the root FortiGate.

To authorize serial numbers of devices from the root FortiGate

When you add the serial number of a Fortinet device to the trusted list on the root FortiGate, the device can join the Security Fabric as soon as it connects. After you authorize the new FortiGate, additional connected FortiAP and FortiSwitch devices automatically appear in the topology tree. From the topology tree, it's easier for you to authorize them with one click.

To authorize a FortiGate or a FortiWiFi from the root FortiGate:

1. Connect to the root FortiGate. To add the serial number of the new FortiGate to the Security Fabric trusted list, enter the following commands:

```
config system csf
  config trusted-list
    edit <serial-number>
  end
end
```

2. To enable FortiTelemetry on an interface, go to **Network > Interfaces** and edit the interface that connects to the FortiGate or FortiWiFi you are authorizing. Under Administrative Access, enable **FortiTelemetry**. For best practices, under Networked Devices, you can also enable **Device Detection**.
3. Connect to the FortiGate you're adding to the Security Fabric and set the following settings on the **Security Fabric > Settings** page:

FortiGate Telemetry	Enable FortiGate Telemetry .
Group name	Set the Group name to the same Security Fabric group name that's configured on the root FortiGate.
Group password	Leave this field blank.
Connect to upstream FortiGate	Enable Connect to upstream FortiGate .
FortiGate IP	Enter the IP address of the root FortiGate or upstream FortiGate you're connecting to.
Apply	Select Apply .

4. Connect to the root FortiGate. Open the **Security Fabric > Settings** page and verify that the FortiGate that you added appears in the Security Fabric Topology.

Joining the Security Fabric by device request

Your device can request to join the Security Fabric from another FortiGate. However, you must have the group name and the IP address of the root FortiGate. The administrator of the root FortiGate in the Security Fabric

must also authorize your device before it can join the Security Fabric.

The root FortiGate must already have FortiTelemetry enabled on the interface that the device is connecting to.

To enable FortiTelemetry on an interface, go to **Network > Interfaces** and edit the interface that connects to the FortiGate or FortiWiFi you are authorizing. Under Administrative Access, enable **FortiTelemetry**. For best practices, under Networked Devices, you can also enable **Device Detection**.

To join the Security Fabric by device request - GUI:

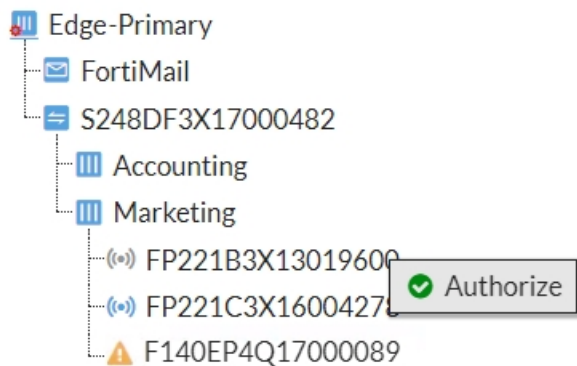
1. Connect to the unauthorized FortiGate or FortiWiFi, and go to **Security Fabric > Settings**.
2. From the Security Fabric Settings page, enable **FortiGate Telemetry**.
3. Enter the group name in the **Group name** field.
4. Leave the **Group password** blank.
5. To connect, enable **Connect to upstream FortiGate**.
6. Set the **FortiGate IP** to the IP address of the root FortiGate or upstream FortiGate that you want to connect to, and select **Apply**.
7. Connect to the root FortiGate and verify that the unauthorized FortiGate appears in the topology tree in **Security Fabric > Settings**. Hover over the unauthorized FortiGate and the tool tip shows the **Status** as **Waiting for Authorization**.
8. To authorize, click on the unauthorized FortiGate and select **Authorize**.

You can also allow other Fortinet devices to join the Security Fabric. You can authorize both FortiAP and FortiSwitch in the Security Fabric with one click. When you connect a FortiAP or FortiSwitch to an authorized FortiGate or FortiWiFi, the device automatically appears in the topology tree.

To authorize FortiAP and FortiSwitch devices

1. The topology tree is in the Security Fabric Settings page and in the Security Fabric Status widget on the Dashboard page. From either widget, click on the grayed out device icon to authorize or deauthorize it. Authorized devices turn blue and unauthorized products disappear from the topology tree.
2. Connect to the upstream FortiGate that the FortiAP or FortiSwitch is connected to.

The image below shows an unauthorized FortiAP in the topology widget:




Note: You can also deauthorize FortiMail from the topology tree, however you must initially authorize FortiMail in the **Security Fabric > Settings** menu.

The following image shows FortiMail options:



Note: You can't authorize or deauthorize a FortiGate from the topology tree widget. You must disable FortiGate telemetry from the FortiGate you wish to deauthorize, or set the serial number to deny on the root FortiGate to remove it from the Security Fabric topology tree.

The following image shows FortiGate options:



Deauthorizing a device

You can deauthorize a device to remove it from the topology tree widget in the Security Fabric Settings page and in the Security Fabric Dashboard.

To deauthorize a FortiGate or FortiWiFi from the root FortiGate - GUI:

1. Connect to the root FortiGate.
2. To deauthorize the serial number of a trusted FortiGate or FortiWiFi, enter the following CLI commands:

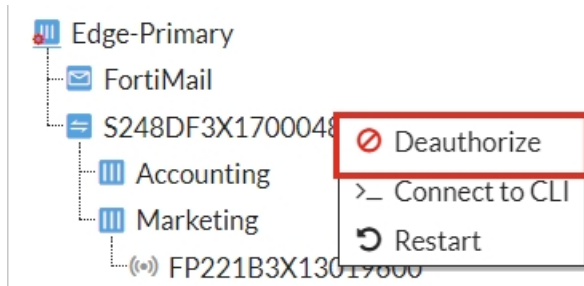
```
config system csf
...
config trusted-list
    edit <serial-number>
        set action deny
    end
end
```

To leave the Security Fabric from a downstream FortiGate or FortiWiFi

1. Connect to the FortiGate or FortiWiFi that you want to deauthorize, and go to **Security Fabric > Settings**.
2. Disable **FortiGate Telemetry**.
3. **Apply** your changes.

To deauthorize a FortiSwitch, FortiAP, or FortiMail - GUI:

1. Connect to the upstream FortiGate and go to **Security Fabric > Settings** to see the topology. Alternatively, you can use the Security Fabric topology widget located in **Dashboard > Main**.
2. Click on the device and select **Deauthorize**. This removes the device from the topology tree.



After deauthorization, the serial numbers of the rejected device are saved in a trusted list that's available only in the CLI. You can view the trusted list using the `show system csf` command. The following example shows how the deauthorized FortiSwitch (from the image above) appears in the `trusted-list` with the action set to deny.

Syntax

```
show system csf
config system csf
  set status enable
  set group-name "Office-Security-Fabric"
  set group-password ENC 1Z2X345V678
  config trusted-list
    edit "FGT6HD391806070"
    next
    edit "S248DF3X17000482"
      set action deny
    next
  end
end
```

Add FortiAnalyzer to the root FortiGate of the Security Fabric

1. In the root FortiGate GUI, select **Security Fabric > Settings**.
2. In the Security Fabric Settings page, enable **FortiGate Telemetry**.
3. **FortiAnalyzer Logging** is automatically enabled.
4. In the **IP address** field, enter the IP address of the FortiAnalyzer that you want the Security Fabric to send logs to. If you select **Test Connectivity**, and this is the first time that you are connecting the FortiGate to the FortiAnalyzer, you will receive an error message because the FortiGate has not yet been authorized on the FortiAnalyzer. You can configure this authorization when you configure the FortiAnalyzer.
5. In the **Upload option** field, select the option for how often you want the FortiGate to send logs to the FortiAnalyzer.
6. If you want log transmissions encrypted, enable the **Encrypt log transmission** option. The log transmissions are encrypted using SSL.
7. Select **Apply**.

Additional CLI commands

You can use the following diagnose commands to view pending authorization requests, accept or deny authorization requests, or troubleshoot commands.

To view pending authorization requests on the root FortiGate - CLI:

```
diagnose system csf authorization pending-list
```

To accept or deny authorization requests to join the Security Fabric - CLI:

```
diagnose system csfd authorization {accept | deny} <serial-number-value>
```

where `serial-number-value` is the serial number of the device that has sent an authorization request to join the Security Fabric.

To view downstream device information - CLI:

```
diagnose system csf downstream
```

Configure ISFW FortiGate devices for the Security Fabric

You must have the group password of the root FortiGate to configure ISFW FortiGate devices for the Security Fabric.

1. In the ISFW FortiGate GUI, select **Security Fabric > Settings**.
2. In the Security Fabric Settings page, enable **FortiGate Telemetry**.
3. In the **Group name** field, enter the group name that you set for the Security Fabric.
4. In the **Group password** field, enter the group password that you set for the Security Fabric.
5. Enable the **Connect to upstream FortiGate** option.
6. In the **FortiGate IP** field, enter the IP address of the port on the upstream FortiGate that this FortiGate connects to. Depending on your network topology, the upstream FortiGate is another ISFW FortiGate or the root FortiGate. The FortiAnalyzer setting is automatically enabled. Settings for the FortiAnalyzer will be retrieved when the ISFW FortiGate connects to the root FortiGate.
7. Select **Apply**.
8. Repeat this procedure on every ISFW in the Security Fabric.

Setting up data collection with FortiAnalyzer

To set up data collection for the Security Fabric, you enable device detection on ISFW FortiGate devices and then connect the FortiAnalyzer to the Security Fabric.

You enable device detection on the interfaces of the ISFW FortiGate devices where you want the devices attached to those interfaces added to the Security Fabric. Only devices detected on those interfaces are shown in the Security Fabric topology views.

Connecting the FortiAnalyzer to the Security Fabric allows the Security Fabric to show historical data for the Security Fabric topology and logs for the entire Security Fabric.

Enable device detection on ISFW FortiGate devices

1. In the ISFW FortiGate GUI, select **Network > Interfaces**.
2. Select the interface that you want to enable device detection on.
3. Select **Edit** and in the **Networked Devices** section, enable **Device Detection**.
4. Select **OK**.
5. Repeat this procedure for every interface that you want to enable device detection on.

Desynchronizing the FortiAnalyzer, FortiSandbox, and FortiManager

If you want to add devices manually, you can edit the **Source IP** for downstream FortiGate devices in the **Central Management** settings. The **Central Management** settings are located in **Security Fabric > Settings**.

However, if you change the **Source IP**, you must change the log settings to `local`.

If you don't want to automatically synchronize the configurations for FortiAnalyzer, FortiSandbox, and FortiManager, you can change the default system settings of the Security Fabric to use local settings.

To use local system settings - CLI:

```
config system csf
    set configuration-sync local
end
```

Where you set the following variables:

Option	Description
default	Synchronizes the configuration for FortiAnalyzer, FortiSandbox, and Central Management to the root FortiGate.
local	Doesn't synchronize the configuration with the root FortiGate, and you must configure settings individually.

Connect the FortiAnalyzer to the Security Fabric



Ensure that all FortiGate devices in the Security Fabric are registered with the same FortiAnalyzer.

1. In the FortiAnalyzer GUI, select **System Settings > Network**.
2. Select **All Interfaces**.
3. Select the port that connects to the root FortiGate.
4. Select **Edit**.
5. In the **IP Address/Netmask** field, enter the IP address used for the Security Fabric configuration on the root FortiGate.
6. In the **Default Gateway** field, enter the IP address of the interface on the root FortiGate that the FortiAnalyzer connects to.
7. Select **OK**.
8. Select **System Settings > Device Manager**.
The FortiGate devices are listed as **Unregistered**.
9. Select the root FortiGate and the ISFW FortiGate devices in the Security Fabric.
10. Select **+ Add Device**.
The FortiGate devices are now listed as **Registered**.
A warning icon will appear beside the root FortiGate, because the FortiAnalyzer requires administrative access to the root FortiGate in the Security Fabric.
11. In the **Authentication** window, complete the **Admin User** and **Password** fields to authenticate the Security Fabric.
After the FortiAnalyzer authenticates the Security Fabric, the FortiAnalyzer shows the full Security Fabric topology.

You can verify that the FortiAnalyzer configuration is successful by selecting **Security Fabric > Settings** on the root and ISFW FortiGate devices. The **Storage usage** field in the **FortiAnalyzer Logging** section should now show storage usage information.



It is recommended that you create a user account for the FortiAnalyzer.

Adding a FortiSandbox to the Security Fabric

The Security Fabric supports both FortiSandbox Appliance and FortiSandbox Cloud. To use FortiSandbox Cloud, you must first activate a FortiCloud account.

To use FortiSandbox in a Security Fabric, you connect the FortiSandbox to the Security Fabric and then configure an antivirus profile to send files to the FortiSandbox. You can also use sandbox inspection in web filtering and FortiClient compliance profiles.

Connect the FortiSandbox to the Security Fabric

You configure FortiSandbox settings on the root FortiGate in the Security Fabric. After you configure these settings, the root FortiGate pushes them to the other FortiGate devices in the Security Fabric.

1. On the root FortiGate, go to **Security Fabric > Settings**.
2. Enable **Sandbox Inspection**.
3. Select either **FortiSandbox Appliance** or **FortiSandbox Cloud**.
4. If you're using a FortiSandbox Appliance, set **Server** to the IP address for the FortiSandbox.
5. Select **Apply**.

To authorize the FortiSandbox appliance, configure the following:

1. On the FortiSandbox, go to **Scan Input > Device**.
2. Edit the root FortiGate.
3. Under **Permissions & Policies**, select **Authorized**.
4. Select **OK**.
5. Authorize the other FortiGate devices in the Security Fabric.

Configure antivirus profiles

1. Go to **Security Profiles > AntiVirus**.
2. Create a new profile, edit an existing profile, or clone and edit an existing profile.
3. Under **Inspection Options**, set **Send Files to FortiSandbox Appliance/Cloud for Inspection** to **All Supported Files**.
4. Enable **Use FortiSandbox Database**.
5. Select **OK**.

Configure web filter profiles

1. Go to **Security Profiles > Web Filter**.
2. Create a new profile, edit an existing profile, or clone and edit an existing profile.

3. Under **Static URL Filter**, enable **Block malicious URLs discovered by FortiSandbox**.
4. Select **OK**.

Configure FortiClient compliance profiles

1. Go to **Security Profiles > FortiClient Compliance Profiles**.
2. Create a new profile, edit an existing profile, or clone and edit an existing profile.
3. Enable **Security Posture Check**.
4. Enable **Realtime Protection** and **Scan with FortiSandbox**.
5. Select **OK**.

Adding a FortiManager to the Security Fabric

When you add a FortiManager to the root FortiGate in the Security Fabric, it automatically synchronizes with any connected Security Fabric devices that are downstream. To add FortiManager to the Security Fabric, you must configure central management on the root FortiGate. Once you configure these settings, the root FortiGate pushes them to the other FortiGate devices in the Security Fabric. The FortiManager must have Internet access.

The following steps also ensure that the FortiGate can receive antivirus and IPS updates and allow remote management through the FortiManager system or FortiCloud service. The FortiManager device provides remote management of a FortiGate over TCP port 541. You must enable the FortiGate management option so the FortiGate can accept management updates to firmware and FortiGuard services.

Registering a FortiGate ensures that it receives updates to FortiGuard services. It also gives you access to technical support. To register the FortiGate, visit the [Fortinet Support](#) website.

To add a FortiManager to the root FortiGate - GUI:

1. On the root FortiGate, go to **Security Fabric > Settings**.
2. Enable **Central Management**.
3. In the **Type** field, select **FortiManager**.
4. Enter the **IP/Domain Name** for the FortiManager.
5. Select **Apply**.
6. On the FortiManager, go to **Device Manager**. The FortiGate devices in the Security Fabric are listed as **Unregistered Devices**.
7. Select the FortiGate devices, then select **+Add**.
8. Select **OK**.

To configure the FortiGate - CLI:

```
config system central-management
  set type fortimanager
  set fmg {<IP_address> | <FQDN_address>}
end
```

For more information about using FortiManager, see "[Central management with FortiManager](#)" on page 2368.

Adding FortiClient EMS to the Security Fabric

You can configure endpoint control for your Security Fabric using FortiClient Endpoint Management System (EMS).



If you disable the **FortiClient Endpoint Management System (EMS)** option found on the **Security Fabric > Settings** page, it deletes all previously configured EMS server entries.

To configure an EMS Server - GUI:

1. To enable endpoint control, go to **System > Feature Visibility** and under Security Features, enable **Endpoint Control**. The FortiClient Endpoint Management System (EMS) section appears in the **Security Fabric > Settings** page.
2. Go to **Security Fabric > Settings** and enable **FortiClient Endpoint Management System (EMS)**.
3. Select the **+** to add it and enter the following:

Name	Enter the name of the EMS server.
Address	Select the FortiClient EMS address from the drop-down menu or select the + to create a new IP address or hostname.
Serial Number	
REST API Calls	

You can add a maximum of 16 EMS Servers.

4. **Apply** your changes.

To configure endpoint control settings - CLI:

```
config endpoint-control settings
    set forticlient-ems-rest-api-call-timeout <value>
end
```

where the value is set between 500 to 30000 milliseconds (default of 5000).

To configure a FortiClient Enterprise Management server - CLI:

```
config endpoint-control forticlient-ems
    edit 1
        set address <firewall-address-name>
        set serial-number <FortiClient-EMS-serial-number>
        set listen-port <listen-port-number>
        set upload-port <upload-port-number>
        set rest-api-auth <FortiClient-EMS-REST-API-authentication>
    next
end
```

where the following values are set to:

Variable	Description
listen-port-number	Set the listening port between 1 and 65535. The default port is 8013.
upload-port-number	Set the uploading port between 1 and 65535. The default port is 8014.

To configure FortiClient registration synchronization settings - CLI:

```
config endpoint-control forticlient-registration-synch
  edit <default-name>
    config {forticlient-winmac-setting | forticlient-android-settings | forticlient-ios-
      settings}
  next
end
```

To configure FortiClient endpoint control profiles - CLI:

```
config endpoint-control profile
  edit <profile-name>
    config {forticlient-winmac-setting | forticlient-android-settings | forticlient-ios-
      settings}
    set forticlient-ems-entries <FortiClient-EMS-entry-name>
  next
end
```

For information about further information about FortiClient EMS, see the [FortiClient EMS Administration Guide](#).

Troubleshooting

The following commands can be useful for testing FortiClient EMS settings, including: signing in or out of FortiClient EMS, quarantining clients using EMS REST API, and adding quarantine calls to the queue. For additional troubleshooting commands, see the [FortiOS CLI Reference](#).

- `diagnose endpoint forticlient-ems-rest-api signin <FortiClient-EMS-entry-name>`
- `diagnose endpoint forticlient-ems-rest-api signout <FortiClient-EMS-entry-name>`
- `diagnose endpoint forticlient-ems-rest-api quarantine-by-ipv4 <ipv4> <FortiClient-EMS-entry-name>`
- `diagnose endpoint forticlient-ems-rest-api unquarantine-by-ipv4 <ipv4> <FortiClient-EMS-entry-name>`
- `diagnose endpoint forticlient-ems-rest-api queue-quarantine-ipv4 <ipv4-address> [,<ipv4-address>...] To add multiple entries, separate the entries by a comma (no spaces).`
- `diagnose endpoint forticlient-ems-rest-api queue-unquarantine-ipv4 <ipv4> [,<ipv4-address>...] To add multiple entries, separate the entries by a comma (no spaces).`
- `diagnose debug application fcnacd_ems <integer>`


Using the Fortinet Security Fabric

Once you set up the Security Fabric, there are various Security Fabric features that you can use to improve your network security, including the following:

- [Understanding the Security Fabric dashboard widgets](#)
- [Viewing the Security Fabric topology](#)
- [Running a Security Fabric Rating](#)
- [Automation stitches](#)

Understanding the Security Fabric dashboard widgets

You can add Security Fabric widgets to the dashboard on the root FortiGate. There are three widgets available: the Security Fabric Status widget, the Security Rating widget, and the FortiMail Statistics widget. The widgets allow you to see information about the status of the Security Fabric when you first log in to the FortiGate.

If any of these widgets do not appear on your dashboard, you can add them using the  settings button in the bottom right corner. On the root FortiGate, select **Dashboard > Main** and the settings button appears when your

mouse hovers over any part of the dashboard.  Select **Add Widget** and under **Security Fabric**, click to add the widget.

The Security Fabric Status widget

The Security Fabric Status widget shows a visual summary of many of the devices in the Security Fabric. You can hover over the icons at the top of the widget to get a quick view of the status of the Security Fabric, including the status of FortiTelemetry and devices in the Security Fabric. You can click to authorize FortiAP and FortiSwitch devices that are connected to an authorized FortiGate.




The widget shows the following information:

- The name of your Security Fabric
- Icons indicating the other Fortinet devices that can be used in the Security Fabric:
 - Devices in blue are connected in your network
 - Devices in gray are unauthorized devices that are connected in your network.
 - Devices in red are not detected in your network, but are recommended for the Security Fabric.
 - An attention icon shows a FortiGate or FortiWiFi waiting for Authorization.
- The names of the FortiGate devices in the Security Fabric

The Security Rating widget


The Security Rating widget shows the latest Security Rating for your Security Fabric. You can configure the widget to show either how your organization's Security Fabric rating compares to the ratings of other organizations that belong to the same industry as your organization or all industries. Your organization's industry is determined from your FortiCare account. You can also configure the widget to either show scores that are specific to your organization's region or all regions. The widget shows the Security Rating score by percentile. To receive a Security Rating score, all FortiGate devices in the Security Fabric must have a valid Security Rating License.



FortiMail Stats widget

The FortiMail Stats widget shows mail detection statistics from FortiMail. You can configure the widget to show statistics from a FortiMail in your Security Fabric. If you have more than one FortiMail in your Security Fabric, you can add additional FortiMail Stats widgets to the Dashboard.






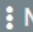


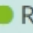



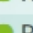
This widget shows both the total number and percentage of email messages that the FortiMail identifies as belonging to non-spam, spam, and virus categories. You can filter the statistics by time period, such as the number of messages per year, per month, and per hour.

FortiMail Stats -  FortiMail		Total ▾	⋮
Total		100	
Non-spam		92 (92%)	
Spam		7 (7%)	
Virus		1 (1%)	


Viewing the Security Fabric topology

You can see the Security Fabric topology in the FortiGate GUI, in the Security Fabric menu. You can choose the Physical Topology or Logical Topology views. In both topology views, you can hover over device icons and use filtering and sorting options to see more information about devices and your organization's network. To view the complete network, you must access the topology views on the root FortiGate in the Security Fabric.

You can also see the Security Fabric topology in the FortiAnalyzer GUI. In the FortiAnalyzer GUI, select **Device Manager**. The FortiGate devices in the Security Fabric are shown as part of a Security Fabric group. An asterisk (*) appears beside the root FortiGate in the Security Fabric. To see the topology of the Security Fabric, right-click on the Security Fabric group and select **Fabric Topology**. Only Fortinet devices are shown in the Security Fabric topology views. The following image shows the FortiAnalyzer GUI.


Device Manager ▾				
	4 Devices Total		1 Devices Unregistered	 0 Devices Log Status Down
+ Add Device  Edit  Delete  More ▾  Column Settings ▾				
<input type="checkbox"/>	▲ Device Name	IP Address	Platform	Logs
<input type="checkbox"/>	✕ Office-Security-Fabric			
<input type="checkbox"/>	Accounting	192.168.65.2	FortiGate-140E-POE	  Real Time
<input type="checkbox"/>	Edge*	192.168.65.2	FortiGate-600D	  Real Time
<input type="checkbox"/>	Marketing	192.168.65.2	FortiGate-81E-POE	  Real Time

The WAN Cloud icon



The WAN cloud icon, in the Physical and Logical Topology views, allows you to receive destination data from the following options in the drop-down menu: Internet, owner IP address, and country/region. These options are available only in the Physical Topology view, when you select **Device Traffic** in the menu in the top right corner.

When you set the WAN cloud icon to **Owner**, the destination hosts are simplified to a fixed size donut chart. This chart shows the percentage division between Internal hosts (with private IP addresses) and Internet hosts. To see which color represents each host, hover over either color. To zoom in on the total number of hosts, click on the donut graph. To see more data for owner details in **FortiView > Destinations**, right-click and select **Destination Owner Details**. You can see the Internet Hosts in the screen shot below.

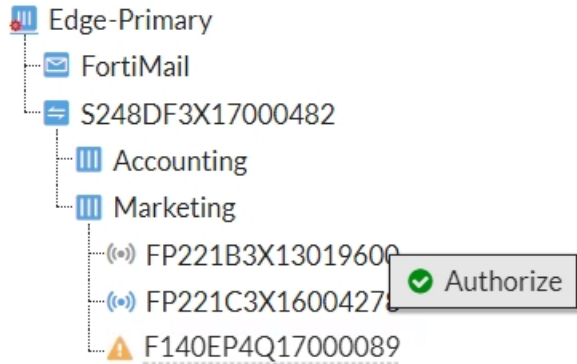


Switch stacking


FortiAP and FortiSwitch links are enhanced in the Security Fabric's Logical and Topological views to show Link Aggregation Groups for the Inter-switch Link (ISL-LAG). This makes it easier to identify which links are physical links and which links are ISL-LAG. To quickly understand connectivity when you look at multiple link connections, ISL-LAG is identified with a thicker single line. To identify ISL-LAG groups with more than two links, you can also look at the port endpoint circles as references.

FortiAP and FortiSwitch integrations

You can see newly discovered FortiAP devices and FortiSwitches in the Security Fabric Topology widget as grayed-out icons, and you can click on any discovered device to authorize it. Once it's authorized, the device icon changes to blue. The following image shows the Security Fabric Topology widget.




The following image shows the same device in the Physical and Logical Topology views.



For an authorized FortiAP, you can right-click to: either **Deauthorize** or **Restart the device**. For an authorized FortiSwitch, you can choose from the following management options: **Deauthorize**, **Connect to CLI**, **Restart** or **Upgrade**.

You can hover over the icon to show the device tooltips. Device tooltips show the connection status and firmware version. FortiSwitch also includes the faceplate and includes the physical port name and number of any devices connected to the FortiSwitch. The following image shows a FortiSwitch device tooltip.



Distinguishing client traffic from server traffic

The Physical and Logical Topology view shows servers and server clusters as rounded square shapes instead of bubbles, and they are grouped separately from endpoint devices to allow you to easily distinguish server traffic from other client traffic. Devices are grouped by device type. For example: Android phones, Apple phones, and Windows PCs.

Identifying compromised hosts from the topology views

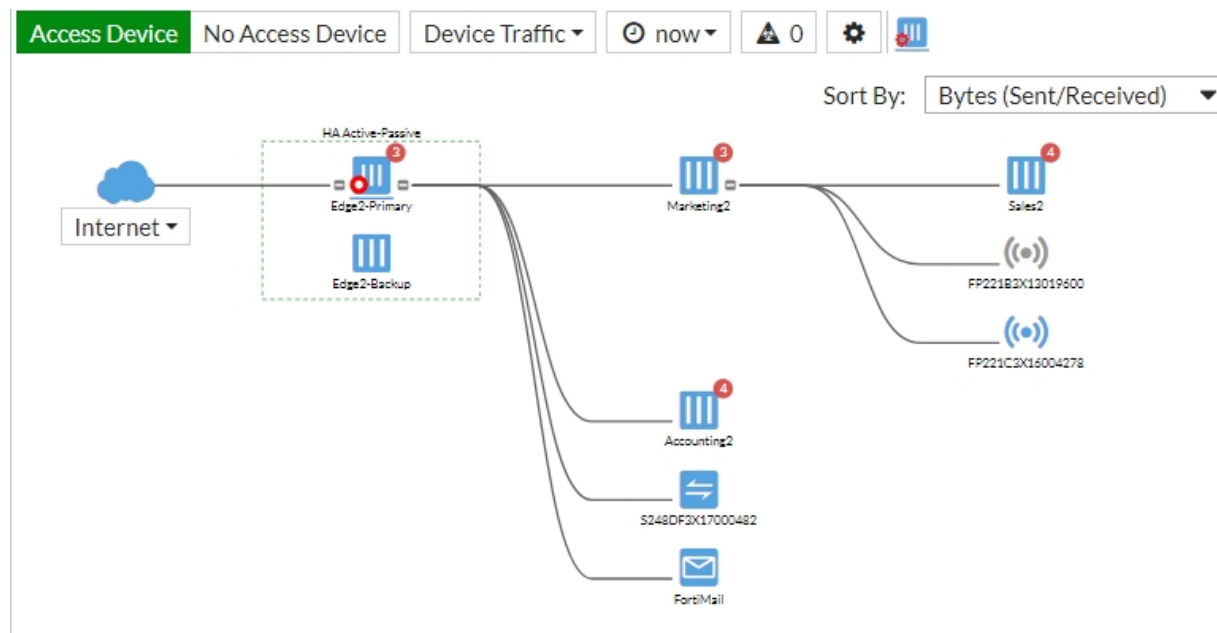
To view compromised hosts on your network and identify threats based on severity, you can click the Compromised Hosts icon in the Physical and Logical Topology views. To filter the compromised hosts by **Device Traffic**, **Device Count** or **IOC Score**, you must set the Bubble Option, located in the top-right corner. The host bubbles show different colors based on the IOC severity level. Confirmed threats appear in an actionable top-down list, located on the right side of the GUI, arranged based on threat severity.

The FortiAnalyzer in the Security Fabric retrieves the detected Indicators of Compromise (IOCs) from FortiGuard services. You can also view compromised host information in the FortiAnalyzer, in **FortiView > Threats > Compromised Hosts**.

View the Physical Topology

The Physical Topology view shows the devices in the Security Fabric and the devices they are connected to. You can also select whether or not to view access layer devices in this topology.

To see the Physical Topology, in the root FortiGate GUI, select **Security Fabric > Physical Topology**.



The Physical Topology view displays your network as a bubble chart of interconnected devices. These devices are grouped based on the upstream device they are connected to. The bubbles appear smaller or larger, based on their traffic volume. You can double-click any bubble to resize it and view more information about the device.

FortiGate devices and other networking devices are depicted as boxes. You can hover over the icon for each FortiGate to see information, such as serial number, hostname, and firmware version. You can hover over the bubbles of other devices to see information about them, such as name, IP address, and traffic volume data.

You can click the **Compromised Hosts** icon , to view compromised hosts on your network and identify threats based on severity.

Security Fabric Rating recommendations are also shown in the topology, beside the icon of the device the recommendations apply to.

View the Logical Topology

The Logical Topology view is similar to the Physical Topology view, but it shows the network interfaces, logical or physical, that are used to connect devices in the Security Fabric.

To see the Logical Topology, in the root FortiGate GUI, select **Security Fabric > Logical Topology**.

The Logical Topology view displays your network as a bubble chart of network connection points. These devices are grouped based on the upstream device interface they are connected to. The bubbles appear smaller or larger, based on their traffic volume. You can double-click any bubble to re-size it.

FortiGate devices and other networking devices are depicted as boxes. You can hover over the icon for each FortiGate to see information, such as serial number, hostname, and firmware version. You can also see each FortiGate interface that has upstream and downstream devices connected to it. You can hover over the name of an interface to see its IP address, network (subnet), and role.

You can click the **Compromised Hosts** icon , to view compromised hosts on your network and identify threats based on severity

Security Fabric Rating recommendations are also shown in the topology, beside the icon of the device the recommendations apply to.

Filter the topology views by specific criteria

You can use filters to narrow down the data on the topology views, so you can find specific information.

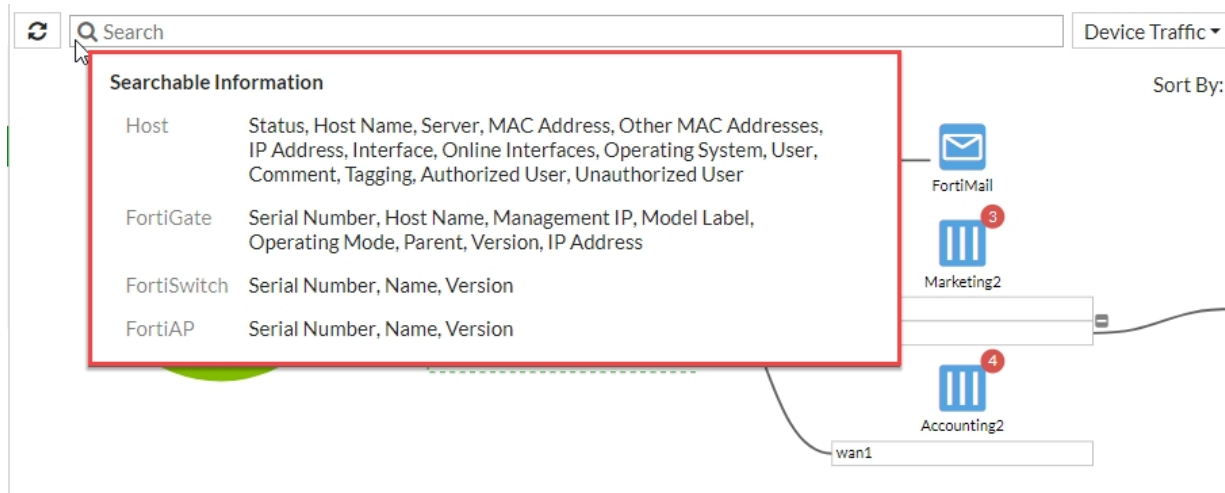
1. In the drop-down menu to the right of the **Search** field, select one of the following:
 - Device Traffic
 - Device Count
 - Device Type
 - Vulnerability
 - Threat Score
 - IOC Score
 - No Device
2. To filter the view by time, in the time period drop-down menu, select one of the following:
 - now
 - 5 minutes
 - 1 hour
 - 24 hours
 - 7 days
3. To sort the topology by traffic options, in the **Sort By** drop-down menu, select one of the following:
 - Bytes (Sent/Received)
 - Packets (Sent/Received)

- Bandwidth
- Session

Using the Search bar to find information in the topology views

The search bar, located above the Physical and Logical Topology views, can help you easily find what you're looking for in the network topology and quickly resolve security issues. For example, you can search for unauthorized hosts and then decide which devices to authorize or remove from your network. The search highlights devices that match your search criteria, and grays out devices that don't match.

To see a list of items that you can search for, mouse over the search bar and a tooltip appears that shows **Searchable Information** list, organized by host and by Fortinet device type. The following image shows the search bar and the **Searchable Information** list:



For hosts, you can search for host information, such as status, host name, and server.

For FortiGate, you can search for device information, such as serial number, host name, and management IP address.

For FortiSwitch and FortiAP, you can search for device information, such as serial number, name and OS version.

Running a Security Fabric Rating

You can run a Security Fabric Rating to analyze your organization's Security Fabric deployment, identify potential vulnerabilities, and highlight best practices that you can use to improve the overall security and performance of your organization's network.

The Security Fabric Rating performs a variety of checks when it analyzes your network. All checks are based on your current network configuration, using real-time monitoring. The check runs across all FortiGate devices in the Security Fabric.

When the check is complete, a list of recommendations is shown. Two views are available: Failed or All Results. You can filter these views further in order to view results from a specific FortiGate or all FortiGate devices. Each view has a chart that shows the results of individual checks, and includes the name and a description of the check, which FortiGate the check was performed on, the impact of the check on the overall security score, and recommendations. If you hover over the result for a check, you can see a breakdown of how the score was determined.

You can choose to automatically apply the recommendations that include the Easy Apply option. By using Easy Apply, you can change the configuration of any FortiGate in the Security Fabric. Further action is required if you want to follow other recommendations.

You can also view recommendations for specific devices in the Physical and Logical Topology views in the Security Fabric menu. If a recommendation is available for a device, a circle containing a number appears. The number shows how many recommendations are available. The color of the circle shows the severity of the highest check that failed. The following table shows the severity that each color represents:

Color	Severity
Red	Critical
Orange	High
Yellow	Medium
Blue	Low

For more information about the Security Fabric Rating, and details about each of the checks that are performed, see the [Fortinet Recommended Security Best Practices](#) document.

Run a Security Fabric check

You must run the Security Fabric Rating check on the root FortiGate in the Security Fabric.

The following image shows the GUI:

Security Rating

1 View Results

2 Easy Apply

All FortiGates

Failed 12

All Results 164

Print

Run Now

Security Rating Score: **+513.9**
Ran: 20 minutes 33 seconds ago

Scheduled Run

112 Passed

8 Medium

4 Critical

FGT6HD3916806070

F140EP4Q17000089

FG81EP4Q16002749

FGT51E3U16002482

Issue	FortiGate	Result	Recommendation
Fabric Security Hardening	4	40	
Firmware & Subscriptions	1		
Network Design & Policies	4		
Threat and Vulnerability Management	3		

Easy Apply >

1. In the root FortiGate GUI, select **Security Fabric > Security Rating**. Click **Show Topology** to view all FortiGate devices in the Security Fabric.
2. To run the check, select **Run Now**.
The check will run. When it completes, it shows the following information:
 - The **Security Rating Score** field shows the score for your Security Fabric
 - The page shows the overall count of how many checks passed or failed, with the failed checks divided by severity
 - Information about each failed check, including which FortiGate failed the check, the effect of the check failure on the security score, and recommendations to fix the issue
 - The **Easy Apply** option appears with recommendations that can be automatically applied by the wizard
3. To move to the **Easy Apply** option page, select **Next**.
4. Select all recommendations that you want to implement in the Security Fabric.
5. Select **Apply Recommendations**.



Not all FortiGate models can run the FortiGuard Security Rating Service if they are the root FortiGate in a Security Fabric. For more information, see the [FortiOS 6.0 Release Notes](#).

Security Fabric Rating license

The Security Rating license is a FortiGuard service and you must purchase a license to access to all the latest features. Security audit checks from FortiOS 5.6 will continue to run, but the following new upgrades are available only when you purchase a Security Rating license:

- Receive FortiGuard updates.
- Run Security Rating checks across each licensed device or all FortiGates in the Security Fabric from the Root FortiGate.
- New 6.0 Rating checks
- Submit rating scores to FortiGuard and receive Security Rating scores from FortiGuard for ranking customers by percentile.

For more information, see the [Fortinet Recommended Security Best Practices](#) document.

Opt out of customer ranking service

You can opt out of submitting Security Rating scores to FortiGuard.

If you opt out of from submitting your network's Security Rating scores, you won't be able to see how your organization's scores compare with the scores of other organizations. Instead, an absolute score is shown.

To disable FortiGuard Security Rating result submission - CLI:

```
config system global
    set fortiguard-audit-result-submission disable
end
```

Logging for Security Fabric Rating

To view the results of past Security Fabric Rating checks, go to **Log & Report > Security Rating Events**.

You can also configure an event filter subtype for the Security Fabric Rating. When you run a check, event logs are created on the root FortiGate that summarize the results of the audit and show detailed information for the individual tests.

To configure logging for the Security Fabric Rating, use the following CLI commands:

```
config log eventfilter
    set security-audit enable
end
```

Understanding the Security Fabric Score

When you run a Security Fabric Rating, your organization's Security Fabric receives a Security Fabric Score. The score will be positive or negative, and a higher score represents a more secure network.

The score is based on how many checks your network passes and fails, as well as the severity level of these checks. The following table shows the weight for each severity level:

Severity level	Weight
Critical	50 points
High	25 points
Medium	10 points
Low	5 points

The check awards points when a check passes, using the following formula:

$$+ <Severity Weight> \times <Secure FortiGate Multiplier>$$

where:

- *Severity Weight* is $<Severity level> / <number of FortiGate devices in the Security Fabric>$
- *Secure FortiGate Multiplier* is determined using logarithms and the number of FortiGate devices in the Security Fabric

For example, if you have four FortiGate devices in the Security Fabric, and all of them pass the Compatible Firmware check, the score for each FortiGate is calculated as: $(50/4) \times 1.292 = 16.2$ points.

All FortiGate devices in the Security Fabric must pass the check in order to receive points. If any of the FortiGate devices in the Security Fabric fail a check, any FortiGate devices in the Security Fabric that passed the check are not awarded points. For the FortiGate that failed the test, the score is calculated using the following formula:

$$- <Severity Weight> \times <Count>$$

where:

- *Severity Weight* is $<Severity level>$
- *Count* is the number of times the check failed during the check

For example, if the check finds two critical FortiClient vulnerabilities, the score for that check is calculated as: $-50 \times 2 = -100$ points.

The score is not affected by checks that do not apply to your network. For example, if you do not have any FortiAP devices in the Security Fabric, you will not receive any points for the FortiAP Firmware Versions check.

Automation stitches



Automation stitches can be used for FortiGate devices that are not part of a Security Fabric.

Automation stitches allow you to decrease response times to security events by automating the activities between different device components in the Security Fabric. You can monitor events from any source in the Security Fabric and set up action responses to any destination.







This section includes:


- [Trigger events](#)
- [Response actions](#)
- [Creating automation stitches](#)
- [Configuring an automation, trigger, and action in the CLI](#)
- [Chaining and delaying actions for AWS Lambda and webhook](#)
- [Diagnose commands for automation stitches](#)

Trigger events

You can configure FortiOS to automatically respond to the following trigger events: IOC, event log, reboot, conserve mode, high CPU, license expiry, HA failover, and configuration changes. The following table provides more information about the trigger event list.

Icon	Trigger	Description
	Compromised Host	<p>An Indicator of compromise (IOC) is detected on a host endpoint.</p> <p>If you configure a Compromised Host trigger you also need to set the IOC level threshold to Medium or High. If you set this to Medium, both medium and high threshold attacks trigger an action.</p> <p>The additional Action options are the following: Access Layer Quarantine, Quarantine FortiClient via EMS, and IP Ban.</p>

Icon	Trigger	Description
	Security Rating Summary	A summary is available for a recently run Security Rating.
	Configuration Change	There is a FortiGate configuration change.
	Reboot	A FortiGate reboot occurs.
	License Expiry	<p>A FortiGuard license is expiring.</p> <p>You must select which type of license you want to be notified about if it expires: FortiCare Support, FortiGuard Web Filter, FortiGuard AntiSpam, FortiGuard AntiVirus, FortiGuard IPS, FortiGuard Management Service, and FortiCloud.</p>
	HA Failover	HA failover occurs.
	AV & IPS DB Update	The antivirus and IPS database updates.

Icon	Trigger	Description
	Event Log	A FortiGate log with a specific event ID occurs. If you configure an Event Log trigger you'll also need to enter a Log ID .
CLI only	Conserve Mode	A FortiGate enters conserve mode due to low memory.
CLI only	High CPU	A FortiGate has high CPU usage.



Response actions



There are four main types of alert notifications you can set up to respond to an event trigger: **Email**, **FortiExplorer Notification**, **AWS Lambda**, and **Webhook**. There are also additional response actions for the Compromised Host (IOC): **Access Layer Quarantine**, **Quarantine FortiClient via EMS** and **IP ban**.




It's recommended that you set a **Minimum Interval** for each action. For more information, see ["Avoiding repeat event notifications" on page 2353](#).



Main Alert Notification Actions

Icon	Action	Description
	Email	Use this action to send a custom email notification. You must enter an email address and subject line.
	FortiExplorer Notification	Use this action to send push notifications to FortiExplorer. For the push to be successful, the FortiGate must be registered with FortiExplorer app on the iOS device you want to receive notifications on.

Icon	Action	Description
	AWS Lambda	<p>Use this action to invoke Amazon Web Services (AWS) Lambda.</p> <p>For the API Gateway endpoint, you can manually enter the URL or you can enter the Parameters individually.</p> <p>For URL, you must enter the following variables:</p> <ul style="list-style-type: none"> • Enter the URL. For example, "1a2b3c.execute-api.us-east-1.amazonaws.com/stagename/notification" • For API Key, enter the same API Key that you use for your AWS API Gateway. <p>For Parameters, you must enter the following variables:</p> <ul style="list-style-type: none"> • Set the Region. For example, "us-east-1" • Set the ID to the REST API ID. For example, "1a2b3c" • Set the Path to the resource you configured in your API Gateway. For example, "notification". • Set the Stage to the stage name from your AWS API Gateway. For example, "stagename". • For the API Key, enter the same API key that you configured in your AWS API Gateway.
	Webhook	<p>Use this action to send data to another application using a REST callback.</p> <p>You must enter the following:</p> <ul style="list-style-type: none"> • For Protocol select HTTP or HTTPS. • For Method select POST, PUT, or GET. • Enter the URI. For example, "website.com/notifications" • Set the Port. • For HTTP Body enter the text you want (up to 1023 characters). For example, {"trigger": "reboot"}. • For HTTP Header, enter the Name and Value you want. For example, "x-notification-source" and "Fortinet".

Additional Compromised Host response actions

Icon	Action	Description
	Access Layer Quarantine	Use this action to impose a dynamic quarantine on multiple endpoints based on the access layer.

Icon	Action	Description
	Quarantine FortiClient via EMS	<p>Use this action to use FortiClient EMS to block all traffic from the source addresses flagged as compromised hosts. Quarantined devices are flagged on the Security Fabric Physical and Logical topology views.</p> <p>Go to Monitor > Quarantine Monitor to view quarantined IP addresses. Addresses are automatically removed from the quarantine after a configurable period of time.</p>
	IP Ban	<p>Use this action to block all traffic from the source addresses flagged by the IOC.</p> <p>Go to Monitor > Quarantine Monitor to view banned IP addresses. Banned IP addresses can only be removed from the list by administrator intervention.</p>

Avoiding repeat event notifications

The **Minimum interval** establishes the amount of time, in seconds, before you receive a repeat alert notification about the same event. This helps avoid receiving multiple alerts on your phone every few minutes for the same offense. When the interval has elapsed, a collated report detailing the activities during that time frame will be sent.

For example, if you were configuring an alert for high CPU usage, and you set the Minimum interval to 86400s (1 day) then you receive one alert when the CPU usage went above 90% and you would not get another alert notification for the same event until the next day. When the 86400s (1 day) elapses, you receive a notification with a summary that let's you know how many times the CPU usage exceeded 90% in the past day.

Creating automation stitches

To create an automation, you can set up a trigger event and response actions that cause the FortiOS to respond in a predetermined way. From the root FortiGate, you can set up triggers for event types, such as compromised host, high CPU, and configuration changes. The automation launches actions in response, such as email alerts, FortiExplorer notifications, and webhooks. The **Compromised Host** trigger has additional actions, such as access layer quarantine and quarantine FortiClient via EMS.

To create and test an automation - GUI:

1. Log in to the root FortiGate, and go to **Security Fabric > Automation**. Select **Create New**.
2. Customize the stitch by selecting a **Trigger** event type and the corresponding **Action** that you would like to automate. You can configure multiple actions for the same event trigger.

Enter the following information:

Name	Enter a name for the new automation.
-------------	--------------------------------------

Status	Select Enabled to enable this automation.
FortiGate	From the drop-down menu, select the FortiGate device to apply this automation to or select All FortiGates (default) to apply to all.
Trigger	<p>Select a Trigger from the following event types:</p> <ul style="list-style-type: none"> • Compromised Host <ul style="list-style-type: none"> • Set IOC level threshold to Medium or High. • Event Log <ul style="list-style-type: none"> • Enter a Log ID. • Reboot • Conserve Mode • High CPU • License Expiry <ul style="list-style-type: none"> • Set the Licensetype to one of the following: FortiCare Support, FortiGuard Web Filter, FortiGuard AntiSpam, FortiGuard AntiVirus, FortiGuard IPS, FortiGuard Management Service, or FortiCloud. • HA Failover • Configuration Changes
Action	<p>If the Trigger event you select occurs, an alert is sent using the methods that you select here. Select at least one of the following Action types:</p> <ul style="list-style-type: none"> • Email <ul style="list-style-type: none"> • Email subject: Enter an email subject. • To: Enter at least one email address. Select the plus + icon to add additional email addresses. • FortiExplorer Notification • AWS Lambda • Webhook <p>NOTE: When you set the trigger to Compromised Host, the following Actions are available:</p> <ul style="list-style-type: none"> • Access Layer Quarantine • Quarantine FortiClient via EMS • IP Ban
Minimum interval (seconds)	Enter a minimum time interval, in seconds, during which you won't receive repeated notifications for the same trigger occurrence. When the minimum time interval expires, you'll receive an alert with a compilation report of any events that occurred during the allotted interval period.

3. Select **OK**.
4. To test the new automation, right-click it and select **Test Automation Stitch**.

When an automation stitch is triggered, the FortiGate creates an event log, which you can view by going to **Log & Report > System Events**.

To create and test an automation - CLI:

```
config system automation-stitch
  edit <automation-stitch-name>
    set status {enable | disable}
    set trigger <trigger-name>
    set action <action-name>
    set destination <serial-number>
  next
end

diagnose automation test <automation-stitch-name> <log>
```



You can configure an automation using the `config system automation-stitch` command shown above. For more information about configuring the **Trigger**<trigger-name> and **Action**<action-name> components, see: ["Configuring an automation, trigger, and action in the CLI" on page 2355](#).

Configuring an automation, trigger, and action in the CLI

This section provides instructions for how to create an automation, and expands on the CLI syntax shown in the introduction by explaining further details, including how to create both a trigger and an action.

To enable the Security Fabric - CLI:

```
config system csf
  set status enable
end
```

To create an "automation-stitch" - CLI:

```
config system automation-stitch
  edit <Automation-stitch-name>
    set status {enable | disable}
    set trigger <trigger-name>
    set action <action-name>
    set destination <serial-number>
  next
end
```

Where the following variables are set:

Variable	Description	Default
<code>edit <Automation-stitch-name></code>	Enter the name of the new automation.	No default

Variable	Description	Default
<code>set status {enable disable}</code>	Enter <code>enable</code> to enable the stitch.	Enable
<code>set trigger <trigger-name></code>	Enter a trigger.	No default
<code>set action <action-name></code>	Enter at least one action you want to occur when a trigger event or schedule occurs.	No default
<code>set destination <serial-number></code>	The <code>destination</code> can be set to a list of device serial numbers, separated by spaces or left blank to use all members of the Security Fabric. Automation stitches are only applied to serial numbers listed in the destination.	All FortiGates

To create an "automation-action" - CLI:

```

config system automation-action
  edit <action-name>
    set action-type {email | ios-notification | alert | disable-ssid | quarantine |
      quarantine-forticlient | ban-ip | aws-lambda | webhook}
    set email-to <email-address>
    set email-subject <subject-name>
    set minimum-interval <seconds>
  next
end

```

Where the following variables are set:

Variable	Description	Default
<code>edit <Automation-action-name></code>	Enter the name of the new automation action.	No default
<code>set action-type</code>	Select an action type from the following: email, ios-notification, alert, disable-ssid, quarantine, quarantine FortiClient, ban IP, AWS Lambda, and webhook.	No default
<code>set email-to <email-address></code>	Enter the email address from which you would like to receive alert notifications. You can add multiple emails by selecting the + icon.	No default
<code>set email-subject <subject-name></code>	Enter the email subject which you would like to see on your email notification alerts.	No default
<code>set minimum-interval</code>	Enter a minimal time interval between 0 to 2592000 seconds, during which a repeat offense of an action will be ignored to help avoid repeat alerts.	Default = 0 seconds

To create an "automation-trigger" - CLI:

```

config system automation-trigger
  edit <trigger-name>
    set trigger-type {event-based | scheduled}
    set event-type {ioc | event-log | reboot | low-memory | high-cpu | license-near-
      expiry | ha-failover | config-change}
    set ioc-level {medium | high}
    set logid [1-99999]
    set license-type {forticare-support | fortiguard-webfilter | fortiguard-antispam
      | fortiguard-antivirus | fortiguard-ips | fortiguard-management | forticloud
      | set trigger-frequency}
    set trigger-frequency {hourly | daily | weekly | monthly}
    set trigger-day <1-31>
    set trigger-hour <0-23>
    set trigger-minute <0-60>
  next
end

```

Where the following variables are set:

Variable	Description	Default
edit <automation-trigger-name>	Enter the name of the new trigger.	No default
set event-type	Select the event type from the following: <ul style="list-style-type: none"> • ioc • event-log • reboot • low-memory • high-cpu • license-near-expiry • ha-failover • config-change 	No default
set ioc-level	Set the IOC level to medium or high. Where: <ul style="list-style-type: none"> • medium sends alerts for both medium and high IOC levels. • high only sends alerts for high IOC levels. NOTE: Only available when event-type is set to ioc.	No default
set logid	Log ID to trigger event. Value from NOTE: Only available when event-type is set to event-log.	No default

Variable	Description	Default
<code>set license-type</code>	<p>Select the license type that you would like to be notified of in the event of expiry. The options include:</p> <ul style="list-style-type: none"> • <code>forticare-support</code> (FortiCare support license) • <code>fortiguard-webfilter</code> (FortiGuard web filter license) • <code>fortiguard-antispam</code> (FortiGuard antispam license) • <code>fortiguard-antivirus</code> (FortiGuard AntiVirus license) • <code>fortiguard-ips</code> (FortiGuard IPS license) • <code>fortiguard-management</code> (FortiGuard management service license) • <code>forticloud</code> (FortiCloud license) <p>NOTE: Only available when <code>event-type</code> is set to <code>license-near-expiry</code>.</p>	No default
<code>set trigger-type</code>	Enter the trigger type as either <code>event-based</code> or <code>scheduled</code> .	No default
<code>set trigger-frequency</code>	<p>How often the trigger is run.</p> <p>The options for the scheduled trigger frequency are the following: hourly, daily, weekly, or monthly.</p> <p>NOTE: Only available when <code>trigger-type</code> is set to <code>scheduled</code>.</p>	Daily.
<code>set trigger-day</code>	Enter an integer value from 1 to 31. This is the day within the month to trigger.	No default
<code>set trigger-hour</code>	<p>Enter the hour of the day on which to trigger from 0 to 23.</p> <p>NOTE: Only available when <code>trigger-type</code> is set to <code>scheduled</code>.</p>	1
<code>set trigger-minute</code>	Enter the minute of the hour on which to trigger (0 - 59, 60 to randomize).	No default

Setting up an automation destination

The `config system automation-destination` command allows you to set the type to the primary FortiGate of an HA cluster or a single FortiGate, and both types of endpoint require it to be set to a destination [by serial number]. Then you can add the destination to any automation stitch. For more information on how to configure an HA cluster as the automation destination see the *High Availability Handbook*.

To set an automation destination:

```
config system automation-destination
  edit <name>
    set type {fortigate | ha-cluster}
    set destination <serial_number>
    set ha-group-id <number>
  next
```

Then you can add the destination to any automation stitch:

```
config system automation-stitch
  edit <stitch-name>
    set destination <destination-name>
  end
```

Chaining and delaying actions for AWS Lambda and webhook

For automation stitches that use the action for **AWS Lambda** or **Webhook**, extra options are available to support chaining and delaying these actions. You can configure these options using the CLI.

To enable chaining actions by delaying an action until the previous action is finished, use the command `set required enable`. This option is only available when `action-type` is set to `aws-lambda` or `webhook`, and it's disabled by default.

To delay the execution of the action, use the command `set delay <seconds>`. This option is only available when `action-type` is set to `aws-lambda` or `webhook`, and it's set to 0 by default.

CLI syntax

```
config system automation-action
  edit <name>
    set action-type {aws-lambda | webhook}
    set required {enable | disable}
    set delay <seconds>
  next
end
```

Diagnose commands for automation stitches

Diagnose commands are available for automation stitches, allowing you to do the following:

- [Test an automation stitch](#)
- [Enable and disable log dumping for automation stitches](#)
- [Display settings for every automation stitch](#)
- [Display history for every automation stitch](#)

Test an automation stitch

To test an automation stitch, use the `diagnose automation test <automation-stitch-name> <log>` command.

Example output

```
# diagnose automation test HA-failover
automation test is done. stitch:HA-failover
```

Enable and disable log dumping for automation stitches

To toggle between enabling and disabling log dumping, use the `diagnose test application autod 1` command.

Example output

```
# diagnose test application autod 1
autod log dumping is enabled

# diagnose test application autod 1
autod log dumping is disabled

autod logs dumping summary:
autod dumped total:0 logs, num of logids:0
```

Display settings for every automation stitch

To display settings for all automation stitches, use the `diagnose test application autod 2` command.

Example output

```
# diagnose test application autod 2
csf: enabled root:yes
total stitches activated: 2

stitch: Compromised-IP-Banned
destinations: all
trigger: Compromised-IP-Banned
actions:
Compromised-IP-Banned_ban-ip type:ban-ip interval:0

stitch: HA-failover
destinations: HA-failover_ha-cluster_25;
trigger: HA-failover
actions:
HA-failover_email type:email interval:0
subject: HA Failover
mailto:admin@example.com;
```

Display history for every automation stitch

To display the history for all your automation stitches, use the `diagnose test application autod 3` command.

Example output

```
# diagnose test application autod 3

stitch: Compromised-IP-Banned

local hit: 0 relayed to: 0 relayed from: 0
last trigger:Wed Dec 31 20:00:00 1969
last relay:Wed Dec 31 20:00:00 1969

actions:
Compromised-IP-Banned_ban-ip:
done: 0 relayed to: 0 relayed from: 0
last trigger:Wed Dec 31 20:00:00 1969
last relay:Wed Dec 31 20:00:00 1969

stitch: HA-failover

local hit: 1 relayed to: 1 relayed from: 1
last trigger:Thu May 24 11:35:22 2018
last relay:Thu May 24 11:35:22 2018

actions:
HA-failover_email:
done: 1 relayed to: 1 relayed from: 1
last trigger:Thu May 24 11:35:22 2018
last relay:Thu May 24 11:35:22 2018
```


Fabric Connectors



You can use Fabric Connectors for FortiGate devices that don't belong to a Security Fabric.

There are three types of Fabric Connectors, which allow you to connect your network to external services. The three types are: SDN, SSO/identity, and threat feeds.

This section contains the following information:

- [Available services for Fabric Connectors](#)
- [Configuring Fabric Connectors](#)
- [Verifying Fabric Connector status](#)
- [Fabric Connector resources](#)

Available services for Fabric Connectors

Fabric Connectors support the following services:

SDN connectors

- Amazon Web Services
- Cisco Application Centric Infrastructure
- Google Cloud Platform
- Microsoft Azure
- Nuage Virtualized Services Platform
- Oracle Cloud Infrastructure
- VMware NSX

SSO/identity connectors

- Poll Active Directory server
- RADIUS single sign-on agent
- Fortinet single sign-on agent

Threat feed connectors

- FortiGuard category
- Firewall IP address
- Domain name



If your FortiGate has VDOMs enabled, SDN and threat feed connectors are global settings, while SSO/identity connectors are available per-VDOM.

Configuring Fabric Connectors

The method that you use to configure a Fabric Connector depends on which type of connector you're using:

- [Creating an SDN Connector](#)
- [Creating an SSO Connector](#)
- [Creating a Threat Feed Connector](#)

Creating an SDN Connector



FortiOS doesn't support multiple SDN Connector instances to Amazon Web Services, Google Cloud Platform, Microsoft Azure, and VMware NSX.

Software-Defined Network (SDN) Connectors provide integration and orchestration of Fortinet products with key SDN solutions. You use SDN Connectors to make sure that any changes in your SDN environment are automatically updated in your network.

To create an SDN Connector, you need to do the following:

- [Gather required information](#)
- [Create the Fabric Connector](#)
- [Create a Fabric Connector address](#)
- [Add the address to a firewall policy](#)

For an example of how to configure a Fabric Connector for Microsoft Azure, see [FortiGate SDN Connector for Azure](#).

Gather required information

Before you can create an SDN Connector, you need to know specific information, which differs depending on which service you're using. You can find this information using your account for the specific service.

Service	Required information for the service
Amazon Web Services	<ul style="list-style-type: none">• Access key ID• Secret access key• Region name• VPC ID (optional)
Cisco Application Centric Infrastructure	<ul style="list-style-type: none">• IP address• Port• Username• Password

Service	Required information for the service
Google Cloud Platform	<ul style="list-style-type: none">• Project name• Service account email• Private key
Microsoft Azure	<ul style="list-style-type: none">• Tenant ID• Client ID• Client secret• Subscription ID (optional)• Resource group (optional)
Nuage Virtualized Services Platform	<ul style="list-style-type: none">• IP address• Port• Username• Password
Oracle Cloud Infrastructure	<ul style="list-style-type: none">• User ID• Tenant ID• Compartment ID• Server region• Certificate
VMware NSX	<ul style="list-style-type: none">• IP address or hostname• Username• Password

Create the Fabric Connector

You can create the Fabric Connector using either the GUI or CLI. The CLI commands that are available vary depending on which service you're using.

Creating a Fabric Connector - GUI:

1. To create a new connector, go to **Security Fabric > Fabric Connectors** and select **Create New**.
2. Select the service you're using and enter the required information for that service.
3. Select **OK**.

Creating a Fabric Connector - CLI:

To create a Fabric Connector using the CLI, use the command `config system sdn-connector`. For more information about this command, see the [FortiOS 6.0 CLI Reference](#).

Create a Fabric Connector address

You use a Fabric Connector address for the following:

- As the source or destination address for firewall policies
- To automatically update changes to the addresses in the environment of the service you're using, based on specified filtering conditions
- To automatically apply changes to the firewall policies that use the address, based on specified filtering conditions

Creating a Fabric Connector address - GUI:

1. To create a new address, go to **Policy & Objects > Addresses** and select **Create New > Address**.
2. Set a **Name** for the address.
3. Set **Type** to **Fabric Connector Address** and set **Fabric Connector Type** to the appropriate service.
4. Set a **Filter**. This filter dynamically creates the members of the address. The types of filters that are supported vary depending on which service you're using.
5. Set a specific **Interface** or leave it as the default **any**.
6. Select **OK**.

Creating a Fabric Connector address - CLI:

```
config firewall address
  edit <name>
    set type dynamic
    set comment <comment>
    set visibility enable
    set associated-interface <interface_name>
    set sdn {aci | aws | azure | nsx | nuage | oci}
    set filter <filter>
  next
end
```

Add the address to a firewall policy

You use a Fabric Connector addresses in a firewall policy as either the source or destination address.

Adding the address to a policy - GUI:

1. To create a new policy, go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Set a **Name** for the policy.
3. Set the appropriate **Incoming Interface** and **Outgoing Interface**.
4. Set the Fabric Connector address as either the **Source** or **Destination** address, as appropriate.
5. Set other policy settings, as required.
6. Select **OK**.

Adding the address to a policy - CLI:

```
config firewall policy
  edit 0
    set name <name>
    set srcintf <port_name>
    set dstintf <port_name>
    set srcaddr <firewall_address>
    set dstaddr <firewall_address>
    set action accept
    set schedule <schedule>
```

```
    set service <service>
  next
end
```

Creating an SSO Connector

SSO Connectors integrate single sign-on (SSO) authentication in your network. SSO allows users to enter their credentials once and have those credentials reused when they access other network resources through your FortiGate.

Connectors are available for the following services:

- Poll Active Directory (AD) server
- RADIUS Single Sign-On (RSSO) agent
- Fortinet Single Sign-On (FSSO) agent

For more information about SSO Connectors, see the [Chapter 3 - Authentication](#).

Creating a Threat Feed Connector

Threat Feed Connectors dynamically import an external block list, in the form of a text file containing a list of either addresses or domains, which resides on an HTTP server. You use block lists to deny access to source or destination IP addresses in web filter and DNS filter profiles, SSL inspection exemptions, and as sources or destinations in proxy policies.

You can configure the following types of threat feeds:

- FortiGuard category
- IP address
- Domain name

For more information about Threat Feed Connectors, see "Overriding FortiGuard website categorization" on page 2455.

Verifying Fabric Connector status



You can only verify the status for Fabric Connectors to AWS, Microsoft Azure, OCI, and VMware NSX.

To verify the status of a Fabric Connector, use one of the following commands:

- `diagnose system sdn status` to verify all connectors
- `diagnose system sdn status <connector_name>` to verify a specific connector

After you enter the command, one of four statuses is displayed:

- `connected`: the connector is connected
- `not connected`: the connector isn't connected
- `disabled`: the related connector entry is set to disabled
- `unknown`: verification of the connector isn't supported

Example output

```
# diagnose sys sdn status
SDN Connector Type Status
-----
aci-sdn-connector aci unknown
aws-sdn-connector aws disabled
azure-sdn-connector azure not connected
nsx-sdn-connector nsx connected
```

Central management with FortiManager

This section describes the basics of using FortiManager as an administration tool for multiple FortiGate devices. It describes the key management features you can use to manage a FortiGate in FortiManager. It contains the following sections:

- [Configuring the FortiManager](#)
- [FortiGuard](#)
- [Firmware updates](#)
- [Administrative domains](#)
- [Backing up and restoring configurations](#)
- [FortiManager in backup mode](#)

For more information about FortiManager, see the [FortiManager Administration Guide](#).



For the FortiGate and the FortiManager to connect properly, both devices must have compatible firmware. To find out if your firmware is compatible, see the [FortiOS Release Notes](#) and [FortiManager Release Notes](#).

Configuring the FortiManager



For information about configuring the connection between your FortiGate devices and FortiManager, see ["Adding a FortiManager to the Security Fabric" on page 2334](#).

After you configure the connection between your FortiGate devices and FortiManager, you can configure the following items on the FortiManager:

- [Configuring updates through FortiManager](#)
- [Using global objects](#)
- [Locking the FortiGate GUI](#)
- [SSL connections](#)

Configuring updates through FortiManager

With the FortiManager system, you can monitor and configure multiple FortiGate devices from one location. You can use the FortiManager Device Manager to view FortiGate devices and make the usual configuration updates and changes, without logging in and out of multiple FortiGate devices.

FortiManager allows you to complete the configuration by going to the Device Manager, selecting the FortiGate, and using the same menu structure and pages as you see in the FortiGate GUI. All changes to the FortiGate configuration are stored locally on the FortiManager until you synchronize with the FortiGate.

When a FortiGate is under the control of a FortiManager device, you shouldn't use the FortiGate to change the configuration. When you try to change options, the FortiGate displays a message stating that it's configured through FortiManager and any changes may be reverted.

Central management configuration supports multiple FortiManager addresses, which helps mainly in the case where the FortiGate is behind NAT.

Using global objects

If you maintain several FortiGate devices within a network, many of the policies and configuration elements are the same across your organization. You can use FortiManager global objects to simplify policy configuration for the FortiGate devices in the network so that you don't have to add and edit many of the same policies on each FortiGate device.

A global object is an object that isn't associated with one device or group. Global objects include security policies, a DNS server, VPNs, and IP pools. You can copy the configurations to the FortiManager device database for a selected device or group of devices. You can also import configurations from the FortiManager device database for a selected device and modify the configuration, as required. When you configure or create a global policy object, the interface, prompts, and fields are the same as creating the same object on a FortiGate using the FortiGate GUI.

Locking the FortiGate GUI

When you use the FortiManager to manage multiple FortiGate devices, a local FortiGate is locked and most administrators are prevented from making configuration changes, using the GUI. The `super_admin` can still make changes to the configuration, but this isn't recommended since it may cause conflicts with the FortiManager.

SSL connections

An SSL connection can be configured to encrypt traffic between FortiManager and the FortiGate devices.

Configuring an SSL connection

Use the following CLI commands in the FortiGate CLI to configure the connection:

```
config system central-management
  set status enable
  set enc-algorithm {default* | high | low}
end
```

The default encryption automatically sets high and medium encryption algorithms. Algorithms used for high, medium, and low follows openssl definitions:

- **High** - Key lengths larger than 128 bits, and some cipher suites with 128-bit keys.

Algorithms are: DHE-RSA-AES256-SHA:AES256-SHA: EDH-RSA-DES-CBC3-SHA: DES-CBC3-SHA:DES-CBC3-MD5:DHE-RSA-AES128-SHA:AES128-SHA

- **Medium** - Key strengths of 128 bit encryption.

Algorithms are: RC4-SHA:RC4-MD5:RC4-MD

- **Low** - Key strengths of 64 or 56 bit encryption algorithms but excluding export cipher suites

Algorithms are: EDH-RSA-DES-CBC-SHA; DES-CBC-SHA; DES-CBC-MD5.

Enable / disable logging of SSL connection events

The following commands allow the user to enable or disable logging of SSL connection events. The default is `disable`.

Syntax

```
config system global
    set log-ssl-connection {enable | disable}
end
```

Enabling or disabling static key ciphers (379616)

The following CLI commands under `system global` let you enable or disable static key ciphers in SSL/TLS connections (e.g., AES128-SHA, AES256-SHA, AES128-SHA256, AES256-SHA256). The default is `enable`.

Syntax

```
config system global
    set ssl-static-key-ciphers{enable | disable}
end
```

FortiGuard

The FortiGuard Distribution Network (FDN) provides FortiGuard services for the FortiManager system and its managed devices and FortiClient agents. The FDN is a world-wide network of FortiGuard Distribution Servers (FDS), which update the FortiGuard services on your FortiManager system on a regular basis so that your FortiManager system is protected against the latest threats.

This section contains the following:

- [Setting up FortiGuard](#)
- [Configuring FortiGuard licensing for devices with limited or no connectivity](#)
- [Troubleshooting your FortiGuard connection](#)

Setting up FortiGuard

In FortiGuard Management, you can configure the FortiManager system to act as a local FDS, or use a web proxy server to connect to the FDN. FortiManager systems acting as a local FDS synchronize their FortiGuard service update packages with the FDN, then provide these updates and look up replies to the FortiGate devices in your private network. The local FDS provides a faster connection, which reduces the Internet connection load and the time that's required to apply frequent updates, such as antivirus signatures, to many devices.

The default port that's used for FortiGuard services is UDP/8888.

The FortiManager system includes the following FortiGuard services:

- Antivirus and IPS engines and signatures (including mobile malware)
- Web filtering and email filtering rating databases and lookups (select systems)
- Vulnerability scan and management support for FortiAnalyzer

To view and configure these services, go to **FortiGuard > Settings** on the FortiManager.

FortiManager can also connect to the FortiGuard Distribution Network (FDN) to receive push updates for IPS signatures and antivirus definitions. These updates can then be used to update multiple FortiGate devices throughout your organization. By using the FortiManager as the host for updates, bandwidth use is minimized as updates are downloaded to one source instead of many.

To receive IPS and antivirus updates from FortiManager, indicate an alternate IP address on the FortiGate.

To configure updates from FortiManager - GUI:

1. In the FortiGate GUI, go to **System > FortiGuard**.
2. Under **AntiVirus & IPS Updates**, enable both **Accept push updates** and **Use override push**.
3. Enter the IP address of the FortiManager.
4. Select **Apply**.

The FortiManager can also operate as a local FDS server when it's in a closed network with no Internet connectivity. For more information, see ["Configuring FortiGuard licensing for FortiGate devices with limited or no connectivity" on page 1](#).

Extended database version OIDs for AV and IPS

New extended database version OIDs ensure accurate display of the AntiVirus and IPS databases in use when you go to **System > FortiGuard**.

IPS signatures page

The IPS signatures list page shows which IPS package is currently deployed. You can change the IPS package by hovering over the information icon next to the IPS package name. Text appears that links directly from the IPS signatures list page to the **System > FortiGuard** page on the FortiGate.

The central management FortiGuard server list can include FQDNs

This feature implements support for FQDN, to make it an option for central-management server-list.

To add FQDN as an address type - GUI:

On **System > FortiGuard > Override FortiGuard Servers > Create New / Edit**, an **FQDN** option, is added for **Address Type**.

CLI changes

```
config server-list
  edit 1
    set server-type {update | rating}
    set addr-type {ipv4 | ipv6 | fqdn} <== added fqdn
    set server-address ipv4
    set server-address6 ipv6
    set fqdn FQDN <== added
  end
end
```

Sending malware statistics to FortiGuard

To support following malware trends and making zero-day discoveries, FortiGate devices send encrypted statistics to FortiGuard about IPS, application control, and antivirus events that FortiGuard services running on

the FortiGate detect. FortiGuard uses the statistics collected to achieve a balance between performance and security effectiveness by moving inactive signatures to an extended signature database.

The statistics include some non-personal information that identifies the FortiGate and its country. This information is never shared with external parties. You can choose to disable the sharing of this information by entering the following CLI command:

```
config system global
    set fds-statistics disable
end
```

Configuring FortiGuard licensing for devices with limited or no connectivity

In some high security environments, Internet service from internal FortiGate devices or for the FortiManager is restricted. This section describes how to configure devices with limited or no internet connectivity to receive FortiGuard updates.

Preliminary steps

1. Register the FortiGate. For a physical FortiGate, use the serial number. For a FortiGate virtual machine (VM), use the registration number. To register the FortiGate, visit the [Fortinet Support](#) website.
2. For FortiGate VMs, the registration process creates a unique license file that's available under **Asset > View/Manage Products**. Select the correct device and download the license file.

This section assumes that:

- Internal FortiGate devices can access a local physical FortiManager or FortiManager VM.
- The FortiManager is running firmware version 6.0.0 or later.

After you have completed the following steps, use the following instructions:

- [Configuring a FortiGate without Internet connectivity to access a local FortiManager as FDN](#)
- [Configure FortiManager without Internet connectivity as a local FDN server](#)

Configuring a FortiGate without Internet connectivity to access a local FortiManager as FDN

By default, FortiGate connects to the public FDN to validate its license and download security feature updates, including databases and engines for security feature updates, such as AntiVirus and IPS. You can configure a FortiGate to use a local FortiManager for both license validation and FDN updates.

For a FortiGate that doesn't have Internet access, you must complete the full configuration before you upload the license. When the FortiGate receives a license file (from the GUI or CLI), it immediately attempts to access the public FDN to validate the license. Until the license is validated, you can't log in to the GUI and some CLI commands aren't available, including the commands that you need to define a local FDN server. This makes it very difficult for you to add the necessary commands to point the FortiGate to a local FortiManager to validate the license.

This document describes how to configure a FortiGate for local FDN access, and provides you with a workaround to fix a FortiGate that can't access a public license validation server.

Follow this procedure to configure a FortiGate to use a local FortiManager for FDN access.



If you complete these steps in a different order, the process may fail, and the FortiGate won't be able to validate the license.

In the FortiGate CLI:

1. Configure central management settings:

```
config system central-management
config server-list
edit 1
set server-type update rating
set server-address <fortimanager_ip>
next
end
set include-default-servers disable
end
```

2. Upload the license using TFTP:

```
execute restore vmlicense tftp <filename>.lic <tftp_ip>
```

The FortiGate reboots.

3. Complete the central management configuration:

```
config system central-management
set fmg <fortimanager_ip>
end
```

In the FortiManager GUI:

You must manually add devices to the FortiManager.

As a result of the CLI commands that you entered on the FortiGate, the device is displayed in the FortiManager GUI, in the **Unregistered Devices** list that's located in the **Device Manager** pane for the root ADOM.

When you enable ADOMs, you must assign the device to an ADOM when you register it.

To add devices manually:

1. Confirm that central management is enabled for the device (as above).
2. In FortiManager, select the root ADOM, and go to **Device Manager**.
3. In the tree menu, click **Unregistered Devices**. The content pane displays the unregistered devices.
4. Select the unregistered device or devices, then click **Add**.
The **Add Device** dialog box opens.
5. If ADOMs are enabled, select the ADOM in the **Add the following device(s) to ADOM** list. If ADOMs are disabled, select **root**.
6. Type the login and password for each device.
7. Click **OK** to register the devices.

The devices are added.

Configure FortiManager without Internet connectivity as a local FDN server

The FortiManager can be operated as a local FDS server when it is in a closed network with no Internet connectivity.

Without a connection to a FortiGuard server, update packages and licenses must be manually downloaded from support, and then uploaded to the FortiManager. Through this feature, known as **Closed Network Mode**, the FortiManager can then provide updates and validate licenses for local FortiGate devices without Internet access.

To configure the FortiManager in Closed Network Mode, complete the following tasks:

- [Enable Closed Network Mode](#)
- [Request FortiGate license validation information](#)
- [Download FortiGuard service update files](#)
- [Configure FortiManager in Closed Network Mode](#)

Enable Closed Network Mode

1. From the FortiManager GUI, go to **FortiGuard > Settings** and disable **Enable Communication with FortiGuard Server**.

Or, from the FortiManager CLI, enable Closed Network Mode by disabling FDS access from the public FDN:

```
config fmupdate publicnetwork
  set status disable
end
```



Once in Closed Network Mode, you must manually import FortiManager service packages, updates, and license upgrades.

Request FortiGate license validation information

1. Create a Customer Service ticket with [Fortinet Support](#) under **Assistance > Create Ticket > Customer Service > Submit Ticket**.
2. Enter the serial number. Under **Category**, select **CS Contract/License**.
3. In the **Comment** field, ask for an "entitlement file" for the FortiGate. Provide the serial number and license number available in **Asset > Manage/View Products > <Select product>**.

Example:

Serial Number: FGVM010000024628

License Number: FGVM0035444



As with asset registration, for large numbers of FortiGate devices you can attach a spreadsheet of serial and license numbers for Customer Service. They will provide a single entitlement file that contains validation information for all FortiGate devices in the spreadsheet. All FortiGate devices must be registered under the same account. Devices registered under different accounts cannot be combined into the same entitlement file.

4. You will receive an entitlement file from Customer Service.

Download FortiGuard service update files

1. From [Fortinet Support](#), navigate to **Download > FortiGuard Service Updates**. Download the **Virus Definition**, **Attack Definition**, and **Mobile Malware** files for the appropriate version of FortiGate and FortiOS. These files are named in the form vsigupdate*.pkg and nids*.pkg.
2. Export the FortiGuard Web Filter and Anti-Spam service updates from a FortiManager that has Internet connectivity by entering the following CLI command:

```
execute fmupdate {ftp | scp | tftp} export <type> <remote_file> <ip> <port>
<remote_path> <user> <password>
```

Enter types `url` and `spam`.

Variable	Description
{ftp scp tftp}	Select the file transfer protocol to use: ftp, scp, or tftp.
<type>	Select the type of file to export or import. The following options are available: av-ips, fct-av, url, spam, file-query, license-fgt, license-fct, custom-url, or domp.
<remote_file>	Update manager packet file name on the server or host.
<ip>	Enter the FQDN or the IP address of the server.
<port>	Enter the port to connect to on the remote SCP host. Range: 1 to 65535 .
<remote_path>	Enter the name of the directory of the file to download from the FTP server or SCP host. If the directory name has spaces, use quotes instead.
<user>	Enter the user name to log into the FTP server or SCP host.
<password>	Enter the password to log into the FTP server or SCP host.

Configure FortiManager in Closed Network Mode

Go to **FortiGuard > Settings** to configure FortiManager as a local FDS server and to upload update packages and licenses.

1. Toggle **OFFEnable Communication with FortiGuard Server** to disable communication with the FortiGuard servers.
2. Toggle **ONEnable Antivirus and IPS Service**.
Select software versions for FortiGate, FortiClient, FortiAnalyzer, and FortiMail.
3. Toggle **ONEnable Web Filter Service**.
When uploaded to FortiManager, the Web Filter database is displayed.
4. Toggle **ONEnable Email Filter Service**.
When uploaded to FortiManager, the Email Filter database is displayed.
5. Under **Upload Options for FortiGate/FortiMail**
 - Upload **AntiVirus/IPS Packages**. Browse for the file you downloaded from the Customer Service Support portal on your management computer. Select **OK** to upload the package to FortiManager. Repeat for each file downloaded from the Customer Service Support portal.
 - Upload **Web Filter Database**. Browse for the file you exported from the FortiManager that is connected to the Internet. Select **OK** to upload the package to FortiManager in the closed network mode. As the database can be large, uploading with the CLI is recommended. See the instructions below.

- Upload **Email Filter Database**. Browse for the file you exported from the FortiManager that is connected to the Internet. Select **OK** to upload the package to FortiManager in closed network mode. As the database can be large, uploading with the CLI is recommended. See the instructions below.
 - Select **Service License** to import the FortiGate license. Browse for the entitlement file on your management computer. Select **OK** to upload the package to FortiManager. A license file can be obtained from support by requesting your account entitlement for the device (see ["Configuring FortiGuard licensing for devices with limited or no connectivity" on page 2372](#)).
6. Under **Upload Options for FortiClient**, select **AntiVirus/IPS Packages** to upload the FortiClient AntiVirus/IPS packages. Browse for the file downloaded from the Customer Service & Support portal on your management computer. Select **OK** to upload the package to FortiManager.

Uploading packages with the CLI

You can upload packages and licenses with the CLI. You should use this method when packages are large, such as database packages.

First, disable communications with the FortiGuard server and enable a closed network with the following CLI commands:

```
config fmupdate publicnetwork
  set status disable
end
```

Then, upload an update package or license by loading the package or license file to an FTP, SCP, or TFTP server.

Run the following CLI command:

```
execute fmupdate { ftp | scp | tftp } import < av-ips | fct-av | url | spam | file-query |
  license-fgt | license-fct | custom-url | domp > <remote_file> <ip> <port> <remote_
  path> <user> <password>
```

Troubleshooting your FortiGuard connection

You can use the following commands to determine the state of license validation and FDN service connectivity, and gather information about connectivity failures. For more information about troubleshooting commands, see the [FortiOS CLI Reference](#).

On a FortiGate, use the following commands:

- `get system status`
- `get webfilter status`
- `get system auto-update version`
- `get system auto-update status`

On a FortiGate VM, use the following commands:

- `diagnose hardware sysinfo vm full`
- `diagnose debug vm-print-license`
- `diagnose hardware sysinfo vminfo`

On a FortiManager, use the following commands:

- `diagnose fmupdate vm-license`

Firmware updates

A FortiManager can also perform firmware updates for multiple FortiGate devices which saves you time because you don't have to upgrade each FortiGate individually.

The FortiManager stores local copies of firmware images, when it downloads images from the Fortinet Distribution Network (FDN) or accepts firmware images that you upload from the management computer.

If you use the FortiManager to download firmware images, the FDN first validates device licenses and support contracts and then provides a list of firmware images that are currently available. For devices with valid Fortinet Technical Support contracts, you can download new firmware images from the FDN and the firmware release notes. After firmware images are downloaded, you can either schedule or immediately upgrade or downgrade the firmware for a device or a group of device.

For more information about updating the FortiGate firmware using FortiManager central management, see the [FortiManager Administration Guide](#).

Administrative domains

FortiManager administrative domains allow the super_admin to create groupings of devices for configured administrators to monitor and manage. FortiManager can manage a large number of Fortinet appliances. This allows you to maintain managed devices that are specific to their geographic location or business division. This also includes FortiGate devices with multiple VDOMs.

Each administrator is tied to an administrative domain (ADOM). When the administrator logs in, they see only those devices or VDOMs that are configured for that administrator and ADOM. The one exception is the super_admin account that can see and maintain all administrative domains and the devices within those domains.

Administrative domains aren't enabled by default and only the super_admin can enable and configure the domains.

The maximum number of administrative domains you can add depends on the FortiManager system model. For more information about the maximums for each model, see the [FortiManager Administration Guide](#).

Backing up and restoring configurations

A FortiManager stores configuration files for backup and restore purposes. A FortiManager also allows you to save revisions of configuration files. Configuration backups occur automatically when you log out or the administrator login session expires.

FortiManager also allows you to view differences between different configurations so that you can identify where changes have been made.

FortiManager in backup mode

Running a FortiManager ADOM in backup mode allows you to use the FortiManager as a central database for address and service objects which are common across multiple devices that are connected to that ADOM. When a FortiManager ADOM is in backup mode, the FortiManager administrator is responsible for managing the database but not the FortiGate configurations.

FortiGate administrators are responsible for making changes to FortiGate devices. When changes are made, the FortiGate administrator is notified of any new, updated, or out-of-sync objects. The administrator can import or update these objects, as needed.

When the FortiManager is in backup mode, the FortiGate configuration is synchronized on demand with FortiManager, similar to how it is done in normal mode and keeps the FortiManager and FortiGate closely synchronized.

Backup mode is useful when you have networks in multiple locations, each with their own administrators who require the ability to make changes quickly.

To add a FortiManager in backup mode, you must first configure an ADOM on the FortiManager that is in backup mode. You can then add the FortiGate to that ADOM using the **Central Management** settings, found at **Security Fabric > Settings**. Make sure to set **Mode** to **Backup**.

To access shared objects from the FortiManager, select the Central Management icon in the GUI header and select **View Details**. A menu opens, showing all FortiManager objects, with options to import, update, or delete objects as necessary.

Related resources

Document	Location
Security Fabric documentation	http://docs.fortinet.com/security-fabric/admin-guides
<i>The Security Fabric Cookbook Recipe Collection</i>	http://cookbook.fortinet.com/security-fabric-collection-60/
<i>Security Fabric Upgrade Guide</i>	https://docs.fortinet.com/security-fabric/release-information
<i>Fortinet Communication Ports and Protocols Guide</i>	https://docs.fortinet.com/fortigate/reference
FortiGate documentation	http://docs.fortinet.com/fortigate/admin-guides
FortiAnalyzer documentation	http://docs.fortinet.com/fortianalyzer/admin-guides
FortiAP documentation	http://docs.fortinet.com/fortiap/admin-guides
FortiClient documentation	http://docs.fortinet.com/forticlient/admin-guides
FortiClient EMS documentation	http://docs.fortinet.com/ems/admin-guides
FortiMail documentation	http://docs.fortinet.com/fortimail/admin-guides
FortiManager documentation	http://docs.fortinet.com/fortimanager/admin-guides
FortiSandbox documentation	http://docs.fortinet.com/fortisandbox/admin-guides
FortiSwitch documentation	http://docs.fortinet.com/fortiswitch/admin-guides
FortiWeb documentation	http://docs.fortinet.com/fortiweb/admin-guides

Chapter 21 - Security Profiles

This FortiOS Handbook chapter contains the following sections:

- [What's New in FortiOS 6.0](#) describes the new security profile features in FortiOS 6.0
- Inside FortiOS highlights the features and benefits of key FortiOS 6.0 components. The technical documentation team maintains these documents as part of the Handbook and as standalone documents which are available at Fortinet's Online Help. Inside FortiOS covers the following security profiles topics:
 - [AntiVirus](#)
 - [Application Control](#)
 - [Intrusion Prevention System](#)
 - [Web Filtering](#)
- [Security profiles overview](#) describes Security Profiles components and their relation to firewall policies, as well as SSL content scanning and inspection.
- [Inspection Modes](#) discusses the FortiGate's inspection modes and how the security profiles function depending on inspection mode.
- [AntiVirus](#) explains how the FortiGate unit scans files for viruses and describes how to configure the antivirus options.
- [Web filter](#) describes basic web filtering concepts, FortiGuard Web Filtering, the order in which the FortiGate unit performs web filtering, and configuration.
- [DNS filter](#) explains how to configure the Domain Name System (DNS) Filter security profile independent of the Web Filter security profile.
- [Application control](#) describes how your FortiGate unit can detect and take action against network traffic based on the application generating the traffic.
- [Intrusion prevention](#) explains basic Intrusion Protection System (IPS) concepts and how to configure IPS options; includes guidance and a detailed table for creating custom signatures as well as several examples.
- [Anti-spam filter](#) explains how the FortiGate unit filters email, how to configure the filtering options, and which actions to take when spam is detected.
- [Data leak prevention](#) describes the DLP features that allow you to prevent sensitive data from leaving your network and explains how to configure the DLP rules, compound rules, and sensors.
- [ICAP support](#) describes how to offload traffic to a separate server specifically set up for the specialized processing of the traffic.
- [FortiClient Compliance Profiles](#) addresses the FortiClient Profiles endpoint protection features and configuration.
- [SSL/SSH Inspection](#) presents SSI and SSH content scanning and inspection with your FortiGate.
- [Custom Application & IPS Signatures](#) describes how to create custom Application Control and IPS signatures.
- [Other security profiles considerations](#) addresses topics like Security Profiles VDOMs, conserve mode, using wildcards and Perl regular expressions, adding External Security Devices, CPU allocation and tuning commands to survive reboot and so on.

What's new in FortiOS 6.0.1

The following list contains new Security Profile features added in FortiOS 6.0.1. Click on a link to navigate to that section for further information.

- ["Option for "Suspicious Files Only" for FortiSandbox submissions" on page 2425](#)

What's new in FortiOS 6.0

The following list contains new Security Profile features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- ["Content Disarm and Reconstruction \(CDR\)" on page 2429](#)
- ["FortiGuard virus outbreak prevention" on page 2432](#)
- ["Overriding the AV engine file scan timeout" on page 2433](#)
- ["Web filtering local and remote category status" on page 2458](#)
- ["Threat Feed Connectors" on page 2462](#)
- ["YouTube Channel Filtering" on page 2465](#)
- ["Configure IPS options" on page 2501 \(for more information, see the \[FortiOS 6.0 CLI Reference\]\(#\)\)](#)
- ["Configuring endpoint protection" on page 2543](#)
- ["Endpoint compliance checking" on page 2545](#)
- ["Enforcing FortiClient EMS requirements" on page 2546](#)
- ["Configuring the FortiClient offline grace period" on page 2548](#)
- ["SSH MITM deep inspection" on page 2564](#)
- ["Global security profiles across Virtual domains \(VDOMs\)" on page 2587](#)
- ["Control how sessions are distributed to Fortinet processes" on page 2592](#)
- ["Excluding industrial IP signatures" on page 2592](#)
- [Extended UTM logging \(for more information, see the \[FortiOS 6.0 CLI Reference\]\(#\)\)](#)

Inside FortiOS: AntiVirus

AntiVirus uses a suite of integrated security technologies to provide against a variety of threats, including both known and unknown malicious codes (Malware), plus Advanced Targeted Attacks (ATA), also known as Advanced Persistent Threats (APT).

Advanced protection against malware and APTs

Malware and Advanced Persistent Threats can cause significant damages to today's organizations. These malicious codes are commonly designed to steal valuable data, gain unauthorized access, or cause products to degrade. FortiOS's AntiVirus is an industry-proven anti-malware security solution with robust features and deployment options

FortiOS offers the unique ability to implement both Flow- and Proxy-based AV concurrently, depending on traffic type, users, and locations. Flow-based AV offers higher throughput performance while proxy-based solutions are useful in mitigating stealthy malicious codes. The AV detection capabilities are further enhanced with complementary security features and external sandbox integration.

By utilizing the unique Content Pattern Recognition Language (CPRL) built into the FortiASIC Content Processor, FortiOS is able to deliver high performance and low latency anti-malware capabilities. This real-time protection is backed by a team of worldwide researchers.

Highlights

- Certification from multiple industries for best-in-class security and capacity with proven coverage and high performance.
- Multi-layered protection with extended AV components and external file analysis integration.
- Comprehensive remediation actions such as file quarantine and knowledge tools.

Key Features & Benefits

Robust feature set	Allows the flexibility to deploy appropriate protection according to security needs and infrastructure designs.
High performance utilizing FortiASIC and patented CPRL AV signatures	Low latency and high capacity ensures that business applications are not affected while security is enforced.
Backed by FortiGuard Labs that deliver real-time protection	Critical digital assets are covered by continuous protection against latest threats.

Features

Industry's validated protection

FortiOS anti-malware components and FortiGuard AV signatures periodically undergo numerous authoritative certifications. These independent certifications demonstrate that the solution offered is of the highest standard in

performance and accuracy, ensuring organizations are truly protected.

Fortinet has been consistently ranked among the top vendors for Virus Bulletin's RAP (Reactive And Proactive) bimonthly tests. This test measures a product's detection rates over the freshest samples available, as well as samples not seen until after product databases are frozen, thus reflecting both the vendor's ability to handle the huge quantity of newly emerging malware and accurately detect previously unknown malware.



Real time protection

The FortiGuard AntiVirus Service provides fully automated updates to ensure protection against the latest content-level threats via the experienced FortiGuard global network is backed by over 200 researchers. With the release of FortiOS 5.6, botnet protection is part of the FortiGuard AntiVirus contract.

FortiGuard AV service quick facts

- 95,000 malware programs neutralized per minute
- 1.8 Million new and updated AV definitions per week
- Hourly updates of the AV signature database
- 190 TB of threat samples till date

Organizations can also engage the FortiGuard Premier Signature Service, which provides enhanced virus detection and threat analysis support. This service offers submissions for custom AntiVirus signatures on a daily basis, offering prioritized support with guaranteed response times. With the release of FortiOS 5.6, botnet protection is part of the FortiGuard AntiVirus contract.

Unique proxy- and flow-based AV

FortiOS offers organizations the flexibility to select the most appropriate inspection method for different network sessions. This can be implemented by defining policies that match specific source objects (IP, IP ranges, users, and devices), destination objects, applications, and schedules with different AV profiles.

Flow-based AV relies on IPS technology where packets are inspected in real-time and matched against the AV signature database. It offers lower latency and higher throughput than Proxy-based AV. Flow-based AV is recommended for inspecting traffic that requires spontaneous user experience or when serving as an additional AV protection layer.

FortiOS's Proxy-based AV offers the most secure AV protection as it's able to inspect more protocols and provides replacement messages on wider range of applications.

AV acceleration with Content Processor

The FortiASICS Content Processor (CP) accelerates content processing traditionally performed completely by the CPU. The CP reduces the resources required by the CPU when matching an incoming file against the signature database, thus improving system performance and stability.

Proactive protection using patented CPRL

Compact Pattern Recognition Language (CPRL) is a patented and proprietary programming language that allows for further inspection of common patterns to not only protect against threats and their variants but also to predict tomorrow's zero-day malware. It allows FortiGuard analysts to describe entire families of malware with a single program, instead of the traditional signature-based "one signature, one variant" model used by other vendors. With fewer signatures to match, throughput performance and latency naturally improve.

Intelligent behavioral evaluation

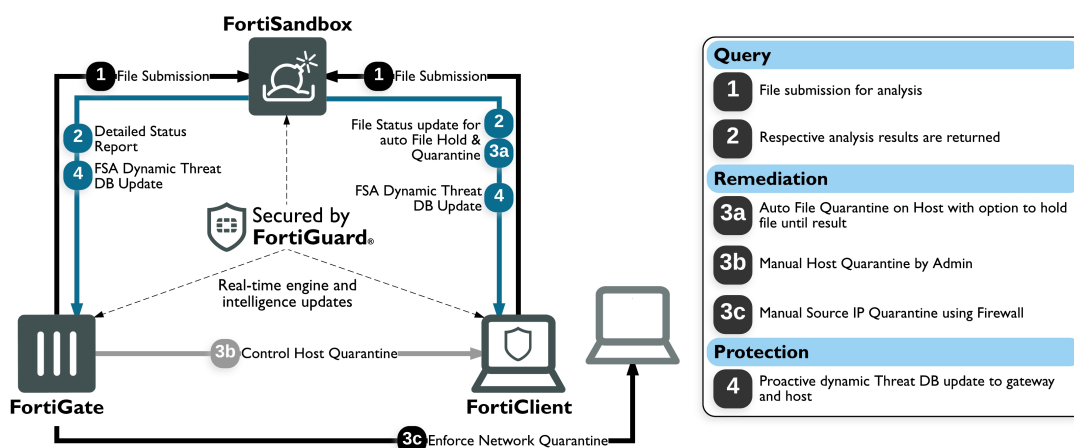
Signature-based security alone is no longer sufficient; it is now critical to understand how devices on your network are behaving. Threat Weight scoring provides a cumulative security ranking of each client device on your network based on a range of behaviors. It provides specific, actionable information that helps identify compromised systems and potential zero-day attacks in real-time.

This unique system attaches predefined scores to various malicious network activities discovered by IPS, application control, URL filtering, etc., to determine the top suspicious users. Administrator can then further inspect these users to undercover unknown threats or APTs via FortiView.

External file analysis integration

FortiOS offers organizations the ability to adopt robust ATP (Advanced Threat Protection) framework that reaches mobile users and branch offices, detecting and preventing advanced attacks that may bypass traditional defenses by examining files from various vectors, including encrypted files. To detect unknown threats, zero-day, and targeted attacks, the FortiGate can engage external resources to perform additional file analysis. Files can be submitted to an on-premise appliance (FortiSandbox) or cloud-based service (FortiSandbox Cloud) after both proxy-based and flow-based AV processing.

It is also possible to configure the FortiGate to automatically receive dynamic signature updates from FortiSandbox and add the originating URL of any malicious file to a blocked URL list. In addition, if the organization deploys integrated endpoint control with FortiClient, an administrator can instruct an infected terminal to self-quarantine.



File filtering

File filtering using data leak prevention (DLP) on the FortiGate offers an effective ways to stop unwanted file transmission instantly. Administrators may implement granular file controls by defining protection profiles using filenames or nearly 50 different file types over mail, web, and file download protocols.

File quarantine

FortiOS offers sophisticated file quarantine capabilities that allow organizations to archive suspicious or blocked files for further examination or to release false positives.

Anti-bot

Organizations may prevent, uncover, and block botnet activities using FortiOS Anti-Bot traffic pattern detection and domain and IP reputation services supplied in real-time by FortiGuard threat experts.

User notification

User notifications are helpful in reducing administration and support burdens, as well as providing user education. FortiOS is able to automatically replace blocked attachments and downloads with detailed information sent to E-mail, FTP, or web users.

Monitoring, logging, and reporting

FortiOS empowers organizations to implement security best practices that require continuous examination of their threat status and adaptation to new requirements. The FortiView widgets provide useful analysis data with detailed and contextual session information, which can be filtered, ranked, and further inspected. System events can also be archived via logs, which in turn can generate useful trending and overview reports.

FortiOS also offers robust in-built E-mail and SMS alert systems, as well as integration with external threat management systems using SNMP and standard-based Syslogs.

Inside FortiOS: Application Control

Application control technologies detect and take action against network traffic based on the application that generated the traffic. Application control uses protocol decoders with signatures that analyze network traffic to detect application traffic, even if the traffic uses nonstandard ports or protocols.

Enhance control and network visibility

Controlling and monitoring applications on a network can seem like a daunting task due to the wide range of available applications. It is no longer an option to simply block or allow TCP and/or UDP ports since most applications do not map to individual ports. For example, controlling traffic on an HTTP or HTTPS port is futile against complex social networking sites and cloud applications.

FortiOS leverages its massive application database to identify applications and their activities while still providing a suitable and sufficient user experience, thanks to FortiASIC Content Processors (CPs), which boost CPU performance. Organizations can adopt more granular control, such as allowing logins but not chatting over selected sites. Traffic shaping may also be applied to the application traffic that is allowed. After applying control measures, continuous monitoring ensures that the measures are effective and allow for changes in application traffic patterns to be managed.

Highlights

- Superior performance using the unique FortiASIC Content Processor that offloads heavy computation from the CPU.
- Flexible implementation with robust deployment modes and granular controls.
- Excellent visibility and management tools that help administrators improve security.
- Application control is a standard part of any FortiCare support contract and the database for Application Control signatures is separate from the IPS database. Access to the database no longer requires a FortiGuard IPS subscription.



Updates for the Application Control signature database require a valid FortiCare support contract.

- Supports detection for traffic using HTTP protocol (versions 1.0, 1.1, and 2.0).
- Ability to configure application control by adding individual applications or application categories to security policies when operating in flow-based inspection and NGFW policy-based mode.

Key features & benefits

Identifies and controls application traffic	Allows organization to strengthen security policies by controlling evasive application communications.
Leverages FortiGate's hardware acceleration and software optimization	Offers more security without compromising performance.

Granular control and integration with other FortiOS capabilities

Provides administrators the ability to implement the most appropriate configuration for any given organization.

Features


NSS Labs “Recommend” rating for Next Generation Firewall

Fortinet’s entry into the NSS Labs Next Generation Firewall Group Test in 2013, 2014 and 2016 received the “Recommend” rating, placing it as one of the top performing systems. NSS Labs uses respectable real-world testing methodologies to measure Next Generation Firewall protection and performance, including application control.

Superior performance with unique hardware architecture

Unlike a traditional security gateway, which relies heavily on CPUs for packet inspection, the FortiGate’s unique hardware architecture allows FortiOS to automatically utilize appropriate hardware components to achieve optimal performance. This prevents the CPU from becoming a bottleneck as it performs various functions concurrently.

In support of application control, the Content Processor (CP) is a specialized ASIC chip that handles demanding cryptographic computation for SSL inspection and intensive signature matching. By offloading these processes from the CPU, the FortiGate is able to minimize performance degradation when administrators opt for greater security.



Robust deployment modes

FortiOS supports a wide array of network protocols and operating modes, allowing administrators to deploy the most appropriate security for their unique IT infrastructure. FortiOS also supports a variety of routing and switching protocols.

The FortiGate is able to operate in inline route and transparent mode. It can also operate in offline sniffer mode for passive monitoring of user activities. These different operating modes run concurrently by using virtual systems.

Protection at the edge

With today’s BYOD and mobile workforce environment, it is no longer wise to deploy control just at the Internet gateway. Through Fortinet Security Fabric, FortiOS unique wireless and switch controller feature allows organizations to implement better visibility and protection closer to internal devices. Moreover, with FortiClient, administrators can also apply similar policies when mobile users are outside of the protected networks.

Advanced application detection and control

By relying on the FortiOS 3rd Generation IPS engine, the FortiGate is able to inspect many of today's encrypted and evasive traffic, as well as traffic running on new technologies, such as SPDY protocol. The inspection can be applied to both network and IPsec/SSL VPN traffic.

An application and its specific activity are identified using FortiGuard's Application Control database of over 2,500 distinct signatures. These signatures are crafted by researchers across the globe to include applications that may be unique to platforms, regions, and/or languages. It also offers specific application activity identification, such as a Facebook posting or Dropbox file sync. The database is kept up to date via scheduled or manual downloads.

The application database is classified into 20 intuitive categories for ease of use. Administrators may also create specific application overrides that differ from the category settings. These specific applications can be filtered and selected by type of behavior, risk levels, technology type, application vendor and popularity.

Administrators may also apply advanced controls, such as setting up session TTLs for specific applications using CLI commands.

Traffic shaping

Organizations may better utilize bandwidth and protect critical applications by enforcing granular application usage with traffic shaping. Administrators can create various traffic shaping profiles by defining traffic priority and maximum or guaranteed bandwidth. These profiles can then be assigned to targeted applications.

User notification

User education is central to an effective security implementation. In response to this, FortiOS lets you provide user notification when blocking an unauthorized application. The notification appears as an HTML block page for web-based applications.

Advanced notification is possible by implementing Fortinet's browser-embedded frame. And when "off-net" users are denied access, notifications appear via FortiClient's notification pop-ups.

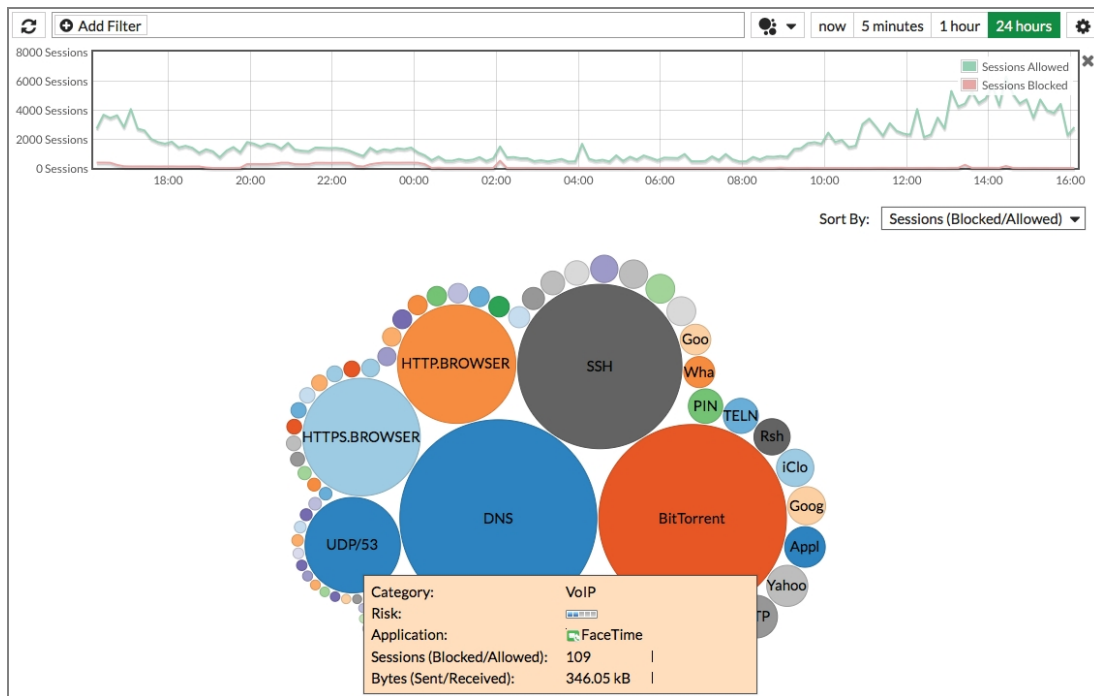
Deep inspection for cloud applications

The prevalence of cloud applications like Dropbox poses a security challenge to today's organizations. Using FortiOS's deep inspection for popular cloud applications, administrators gain deep and useful insights, via FortiView and logs, into activities associated with these applications, such as user IDs, cloud actions, file names, and file sizes. For popular video sites, FortiOS will also be able to track video files viewed.

SSL inspection for encrypted traffic

SSL (Secure Sockets Layer) is a popular encryption standard used to protect Internet traffic but may also be used to evade traditional inspection. FortiOS enables organizations to adopt effective application control even when traffic is encrypted.

Unique hardware components and software optimizations can decrypt traffic with minimal performance impact. The inspection can easily omit sensitive communications, such as financial transaction (thereby complying with privacy policies), or bypass applications that forbid SSL inspection by using granular policy settings.



Monitoring, logging, and reporting

FortiOS empowers organization to implement security best practices that require continuous examination of threat statuses and the ability to adapt to new requirements.

The FortiView widgets provide useful analyses with detailed and contextual session information that can be filtered, ranked, and further inspected. For example, an administrator can instantly query the top applications that are currently consuming bandwidth and drill down to identify their users and help decide if such activities should be blocked.

Network, threat, and system events activities can be archived via syslogs. In turn, these logs can generate useful trending and overview reports.

Lastly, the FortiOS offers robust in-built email and SMS alert systems. Meanwhile, integration with external threat management systems can be achieved with SNMP and standard-based syslogs.

Recipes

Visit cookbook.fortinet.com for these and other recipes:

- NGFW policy-based mode

Inside FortiOS: Intrusion Prevention System (IPS)

Intrusion Prevention System (IPS) technology protects your network from cybercriminal attacks by actively seeking and blocking external threats before they can reach potentially vulnerable network devices.

World class next generation IPS capabilities

Today, sophisticated and high volume attacks are the challenges that every organization must recognize. These attacks are evolving, infiltrating ever-increasing vectors and complex network environments. The result is an urgent need for network protection while maintaining the ability to efficiently provide demanding services and applications.

FortiOS's IPS functionality is an industry-proven network security solution that scales up to over 200 Gbps of in-line protection. Powered by purpose-built hardware and FortiASICs, FortiOS is able to achieve attractive TCO while meeting performance requirements. IPS is easy to set up, yet offers feature-rich capabilities, with contextual visibility and coverage. It is kept up-to-date by research teams that work 24 hours a day worldwide, in order to detect and deter the latest known threats as well as zero-day attacks.

Highlights

- Validated best-in-class security and capacity with proven coverage and high performance.
- Comprehensive protection provided by a signatures-based IPS engine, protocol anomaly scanning, and DDOS mitigation.
- Flexible deployment options and actionable implementations for a wide array of network integration and operation requirements.

Key features & benefits

High Performance IPS, powered by FortiASIC	Low latency and high capacity ensure business applications are not affected while security is enforced.
Best-in-class security with superior coverage	Protects critical digital resources from both internal exploits and external cybercriminals, even if sophisticated attacks are crafted.
Backed by FortiGuard Labs that deliver real-time protection	Maintains up-to-date and proactive protection against latest known threats and newly discovered hacking techniques while allowing time for organizations to patch vulnerable systems.

Features

Tested and proven protection

Not only have FortiGates been deployed in some of the largest enterprises in the world since 2002, FortiOS IPS components and FortiGuard IPS signatures are periodically tested and certified by well-known external labs. For example, Fortinet's FortiGate 3000D earned the highest ratings for Security Effectiveness, blocking 99.9 percent

of exploits in the recent NSS Labs DCIPS test. These independent certifications ensure that solutions delivered to customers are of the highest standards in performance, coverage, and accuracy.

Real-time & zero-day protection

The FortiGuard Intrusion Prevention Service (IPS) provides customers with the latest defenses against stealthy network-level threats through a constantly updated database of known threats and behavior-based signatures.

FortiGuard IPS service quick facts

- Over 10,000 signatures consisting of 18,000 rules
- Approximately 470,000 network intrusion attempts resisted per minute
- About 1,000 rules are updated or added per week
- Over 300 Zero-day vulnerabilities discovered to date

This update service is backed by a team of threat experts and a close relationship with major application vendors. The best-in-class team also uncovers significant zero-day vulnerabilities continuously, providing FortiGate units with advanced protection ahead of vendor patches.

Uncompromised performance

The FortiASICS Content Processor (CP) accelerates content processing, which is traditionally done completely by the CPU. The CP reduces the resources required by the CPU when matching an incoming file against the signature database, thus improving system performance and stability.

Protocol decoders and anomaly detection

Protocol decoders are required to assemble the packets and detect suspicious, nonconforming sessions that resemble known attacks or are non-compliant to RFC or standard implementation.

FortiOS offers one of the most comprehensive arrays of protocol decoders in the industry, providing customers with significantly wide coverage in all kinds of environments.

Pattern & rate-based signatures

The pattern signature matching technique is essential in IPS implementation due to its high level of precision and accuracy. FortiOS offers administrators robust pattern signature selection using filters based on severity, target, operating system, application, and protocol. Each of the 10,000+ signatures has a direct link to its detailed entry on the threat encyclopedia and CVE-ID references. After selection, administrators are able to assign associated actions such as monitoring, blocking, or resetting the session.

Rate-based IPS signatures protect networks against application based DoS and brute force attacks. Administrators can configure nearly 30 rate-based IPS signatures and tune them to their needs. Threshold (incidents per minute) and an action to take when the threshold is reached can be assigned to each signature. If the action is set to block, then a timeout period can be set so that the block is removed after a specified duration.

DoS and DDoS mitigation

DoS policies can help protect against DDoS attacks that aim to overwhelm server resources. In FortiOS, the DoS scans precede the policy engine at the incoming interfaces, thus eliminating unnecessary sessions from the

firewall process and state table entry during a surge of attack traffic. This helps to safeguard the firewall from overloading and allows it to perform optimally.

FortiOS DoS policies can be configured to detect and block floodings, port scans, and sweeps. Administrators can set baselines for the amount of concurrent sessions from sources or to destinations. The settings utilize thresholds and can be applied to UDP, TCP, ICMP, IP, and SCTP.

Network interfaces associated with a port attached to a Network Processor (NP) can be configured to offload anomaly checking, further offloading the CPU for greater performance. Some of the anomaly traffic dropped includes LAND attacks, IP protocol with malformed options, and WinNukes.

Quarantine attacks

FortiOS offers sophisticated automatic attack quarantine capabilities which allow organizations to proactively prevent further attacks from known attackers over a predefined duration. Quarantining by duration can be used to protect potentially vulnerable servers until more permanent defense.

Packet logging

Administrators may choose to automatically perform IPS packet logging, which saves packets for detailed analysis when an IPS signature is matched. Saved packets can be viewed and analyzed on the FortiGate unit or by using third-party analysis tools. Packet logging is also useful in determining false positives.

Custom signatures

Custom IPS signatures can be created to further extend protection. For example, you can use custom IPS signatures to protect unusual or specialized applications, or even custom platforms from known and unknown attacks.

Organizations may use FortiConverter to easily convert Snort signatures for FortiOS use.

Edit IPS Sensor
default
[View IPS Signatures]

Name: default
Comments: Prevent critical attacks. 25/255

IPS Signatures

+ Add Signatures
Delete
Edit IP Exemptions

Name	Exempt IPs	Severity	Target	Service	OS	Action	Packet Logging
No matching entries found							

IPS Filters

+ Add Filter
Edit Filter
Delete

Filter Details	Action	Packet Logging
Severity: Medium Medium, High High, Critical Critical	Default	✖

Rate Based Signatures

Enable	Signature	Threshold	Duration (seconds)	Track By	Action	Block Duration (minutes)
<input checked="" type="checkbox"/>	Apache.HTTP.Server.DoS	200	1	Any	✖ Block	None
<input checked="" type="checkbox"/>	Digium.Asterisk.File.Descriptor.DoS	20	1	Any	✖ Block	None
<input checked="" type="checkbox"/>	Digium.Asterisk.IAX2.Call.Number.DoS	275	1	Any	✖ Block	None
<input checked="" type="checkbox"/>	DotNetNuke.Padding.Oracle.Attack	1000	5	Any	✖ Block	None
<input checked="" type="checkbox"/>	FTP.Login.Brute.Force	200	10	Any	✖ Block	None
<input checked="" type="checkbox"/>	FreeBSD.TCP.Reassembly.DoS	10	2	Any	✖ Block	None
<input checked="" type="checkbox"/>	GlassFish.Login.Brute.Force	200	10	Any	✖ Block	None

Apply

Resistant against evasions

Evasion techniques attempt to fool the protocol decoders in IPS products by crafting exotic network streams that would not be handled or reconstructed by the decoders, yet still be valid enough for the target recipient to process. Robust IPS engine is capable of handling both common evasions and sophisticated AETs (Advanced Evasion Techniques) deployed by hackers such as IP Packet Fragmentation, TCP Stream Segmentation, RPC Fragmentation, URL & HTML Obfuscation, and other protocol specific evasion techniques.

Intrusion detection mode

In out-of-band sniffer mode (or one-arm IPS mode), IPS operates as an Intrusion Detection System (IDS), detecting attacks and reporting them but not taking any action against them. In sniffer mode, the FortiGate unit does not process network traffic and instead is connected to a spanning or mirrored switch port, or a network tap. If an attack is detected, log messages can be recorded and alerts sent to system administrators.

Traffic bypass

Since most IPS deployments are in transparent inline mode, active traffic bypass is often desired until normal operation of the device resumes. Some FortiGates offer inbuilt active bypass interfaces while others may use external bypass devices such as the FortiBridge. Administrators are also offered with software fail-open option to tackle instances where the IPS engine fails.

Monitoring, logging, and reporting

FortiOS empowers organizations to implement security best practices that require continuous examination of their threat status and adaptation to new requirements. The FortiView query widgets provide useful analysis data

with detailed and contextual session information, which can be filtered, ranked, and further inspected. System events can also be archived via logs, which in turn can generate useful trending and overview reports.

Inside FortiOS: Web Filtering

A Web Filtering solution is designed to restrict or control the content a reader is authorized to access, delivered over the Internet via the Web browser. It may be used to improve security, prevent objectionable activities, and increase productive within an organization.

Intelligent and effective content control

Web-based threats such as Phishing, drive-by Malware sites, and Botnets are more sophisticated and scrutinized than ever, and as well as increasingly difficult to control due to the rise of mobility in the workplace, even more difficult for you to control. The Web has become the preferred medium of choice for hackers and thieves looking for new ways to disrupt services, steal information, and perform malicious activities for financial gain. In addition, employees who visit websites containing objectionable content can expose your organization to civil or criminal liability.

FortiOS Web Filtering solution utilizes three main components of the web filtering function: the Web Content Filter, the URL Filter, and the FortiGuard Web Filtering Service. These functions integrate with each other to provide maximum control over what the Internet user can view as well as protection to the network from many Internet content threats. Web Content Filtering blocks web pages containing words or patterns that you specify. URL filtering uses URLs and URL patterns to block or exempt web pages from specific sources. FortiGuard Web Filtering provides many additional categories you can use to filter web traffic by independent real-world tests.

Highlights

- Comprehensive and advanced Web Filtering features Safe Search and user override options.
- FortiGuard Web Filtering Services with superior coverage of over 250 million rated websites.
- Integration with other FortiOS components, such as User Identification for flexible and secured implementation.
- Supports detection for traffic using HTTP protocol (versions 1.0, 1.1, and 2.0).
- Ability to configure web filtering by adding URL categories to security policies when operating in flow-based inspection and NGFW policy-based mode. You can set the action to accept or deny to allow or block the applications.

Key features & benefits

Cloud-based Rating Database	Real-time website category rating provides accurate content control.
Wide choice of web filtering technologies	Various web filtering technology options are available to provide each organization the most suitable implementation.
Integrated with other security and networking functions	Allows organizations to simplified networks and reduce TCO.

Features

Cloud-based rating system

Fortinet is a pioneer in cloud-based rating systems for web filtering. FortiOS provides an innovative approach to HTTP and HTTPS web filtering technology by combining the advantages of a cloud-based service offering with layered response caching. The multiple FortiGuard data centers around the world hold the entire categorized URL database and receive rating requests from FortiGate units triggered by browser-based URL requests. FortiGuard responds to these rating requests with the categories stored for specific URLs, the requesting FortiGate unit then uses its own local profile configuration to determine what action is appropriate to the category, such as: blocking, monitoring, allowing the page, displaying a warning, or requiring authentication to view the page.

Rating responses are also cached directly in FortiGate unit memory so that ratings for frequently used sites can be retrieved directly from the cache, reducing the number of requests to the FortiGuard network. Caching URLs in memory makes URL lookups almost instantaneous while only using a very small amount of system memory.

An appropriately licensed FortiManager appliance can be synchronized to the FortiGuard network and as such can be used in the same way to as the FortiGuard network for managed FortiGate devices. This can further reduce any latency associated with the round trip time for individual rating requests while at the same time ensuring complete database coverage. Consider the combination of a LAN attached FortiGate cluster and FortiManager combination with the potential to handle tens of thousands of requests per second.

Superior coverage

FortiGuard Web Filter ratings are performed by a combination of proprietary methods including text analysis, exploitation of the web structure, and human raters. This service currently rates more than 250 million sites covering billions of URLs with each site able to be rated in multiple categories. The FortiGuard database provides a truly international service with support for 70 languages.

Extensive and flexible categorization

Rated URLs are assigned into one of the 98 categories (including 20 user defined ones) which administrators can then easily manage and control. Administrators can configure and populate local categories or place specific URLs in existing categories should the FortiGuard rating not be in agreement with an organization's policies and practices.

Rating override

At times, administrators may have to allow approved people to access what they need during periods when an exception to the normal rules is required, while still having enough control that the organization's web usage policies are not compromised. FortiOS can provide such setup by using alternate profiles.

Protection against malicious URLs

The malicious URL database contains all malicious URLs active in the last month and is organized as one of the categories. With Fortinet Security Fabric, customers can further their protection by having the FortiSandbox add newly discovered URLs to a dynamic URL filter, thus blocking files from being downloaded again from that URL.

Inspection modes

FortiOS web filtering can operate in different modes: proxy-based and flow-based inspection modes and DNS filtering. Each mode has strengths and weaknesses and all three can be active at the same time on different traffic streams.

Proxy-based web filtering uses a proxy to assemble and analyze web content as it passes through the FortiGate unit. If a page is blocked the proxy can replace the blocked page with a customizable web page informing users that the page is blocked. Proxy-based web filtering is the most feature-rich mode, supporting many advanced filters including web content filtering that analyzes web page content according to your custom requirements, Java applet filtering, and blocking invalid URLs.

Flow-based web filtering uses the FortiOS IPS engine to filter web content packets as they pass through the FortiGate unit without any buffering. Flow-based inspection does not use a proxy, so inspected packets are not proxied and altered by the FortiGate unit. Flow-based inspection does not support as many advanced features as proxy-based web filtering.

To control your FortiGate's security profile inspection mode in FortiOS 5.6, you can select **Flow** or **Proxy Inspection Mode** from **System > Settings**. Having control over flow and proxy mode is helpful if you want to ensure that only flow inspection mode is used.

In most cases proxy mode is preferred because more security profile features are available and more configuration options for these individual features are available. Some implementations, however, may require all security profile scanning to only use flow mode. In this case, you can set your FortiGate to flow mode knowing that proxy mode inspection will not be used.

Two new policy modes are available in FortiOS 5.6.

- NGFW mode simplifies applying application control and web filtering to traffic by allowing you to add applications and web filtering profiles directly to policies. This is used in conjunction with flow-based inspection.
- Transparent proxy allows you to apply web authentication to HTTP traffic without using the explicit proxy.

DNS web filtering employs DNS lookups to the FortiGuard DNS service to get web page ratings. Filtering is done as part of the DNS lookup and web pages can be blocked or redirected to a web filter block page before the HTTP session starts. As a result, it is lightweight in terms of resource usage although it only supports a limited number of advanced features.

Usage quota

Administrators can set a daily timed access quota by category or category group. Quotas allow access for a specified length of time or traffic volume, calculated separately for each user.

SafeSearch

SafeSearch is a feature of popular search sites that prevents explicit web sites and images from appearing in search results. Although SafeSearch is a useful tool, especially in educational environments, the resourceful user may be able to simply turn it off. Enabling SafeSearch on the FortiGate for the supported search sites can better enforce its use by rewriting the search URL to include the code to indicate the use of the SafeSearch feature.

Restrict YouTube access

In FortiOS 5.6 with inspection mode set to proxy-based, you can set Strict or Moderate access to YouTube in a Web Filter profile.

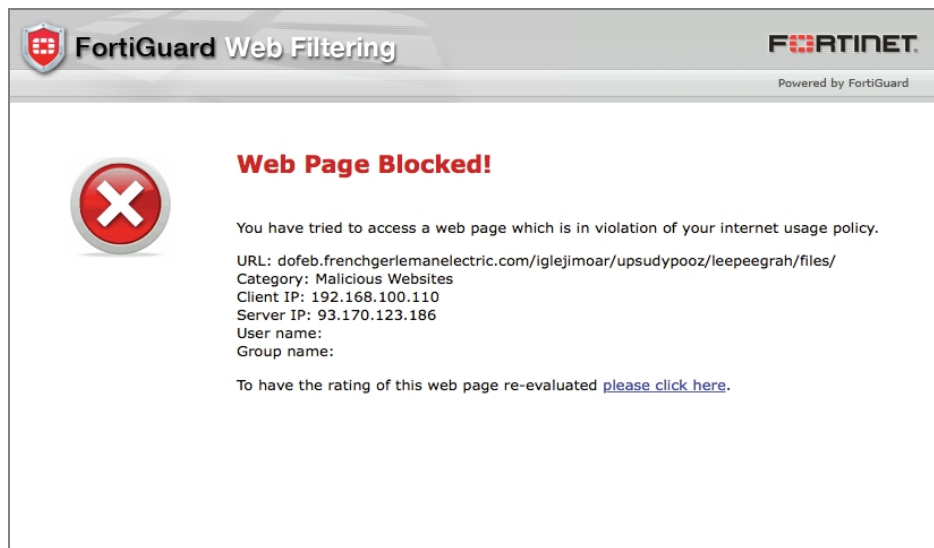
Manual URL and content filter

FortiOS web filtering offers specific URL filtering by standard, wildcard, and regular expression definition, as well as content filtering by pattern type and language.

Advanced web filter configurations

FortiOS rich feature set includes ability to implement a number of enterprise features such as:

- Block HTTP redirects by rating, invalid URLs, HTTP POST actions, and Web resume download
- Cookie, Java applet, and ActiveX filter
- Rate Images by URL and URLs by domain and IP address
- Restrict Google account usage to specific domains



Proxy avoidance preventions

FortiGate is able to improve the effectiveness of the web filtering by preventing users from evading the security implementation. Organizations can use its multiple integrated technologies including proxy site URL, proxy application control, and IPS proxy behavior blocking.

User and device awareness

Most networks in today's organizations are connected with both corporate and personal mobile devices. User and device awareness provides the option to configure intelligent policies that can effectively enforce security.

To tackle the prevalence of BYOD environments, administrators are able to configure web content access policies with sources defined by IPs, users, and devices, either combined or selectively.

External URL filtering support

In instances where customers have large, existing, deployed implementations of a specific URL filtering solution but replace their legacy firewalls with a FortiGate family, they can still retain their web filtering infrastructure since FortiOS supports both ICAP and WISP.

Monitoring, logging, and reporting

FortiOS empowers an organization to implement security best practices that require continuous monitoring of threats, allowing the organization to adapt to new requirements.

The FortiView dashboards display useful analysis data with detailed and contextual session information, which can be filtered and ranked, with drilldown options also available. This information, including system events activities and administration audit trails, can also be archived via logs.

FortiOS logs all the types of traffic that can connect to or terminate at the FortiGate unit. In turn, these logs can generate useful trending and overview reports.

Security profiles overview

The FortiGate line combines a number of security features to protect your network from threats. As a whole, these features, when included in a single Fortinet security appliance, are referred to as Security Profiles.

This overview addresses the following topics:

- [Traffic inspection](#)
- [Content inspection and filtering](#)
- [Security profile components](#)
- [Security profiles/lists/sensors](#)

Firewall policies limit access, and while this and similar features are a vital part of securing your network, they are not covered in this discussion of Security Profiles.



FortiOS 5.4 no longer supports FortiClient 5.0.

FortiOS 5.4.1 supports only FortiClient 5.4.1. Be sure to upgrade managed FortiClients before upgrading the FortiGate to 5.4.1.

FortiOS 5.2 can support FortiClient 5.0, but only if the FortiGate upgraded to FortiOS 5.2. Customers need to purchase a FortiClient 5.4 subscription-based FortiClient license.

Traffic inspection

When the FortiGate unit examines network traffic one packet at a time for IPS signatures, it is performing traffic analysis. This is unlike content analysis where the traffic is buffered until files, email messages, web pages, and other files are assembled and examined as a whole.

DoS policies use traffic analysis by keeping track of the type and quantity of packets, as well as their source and destination addresses.

Application control uses traffic analysis to determine which application generated the packet.

Although traffic inspection doesn't involve taking packets and assembling files they are carrying, the packets themselves can be split into fragments as they pass from network to network. These fragments are reassembled by the FortiGate unit before examination.

No two networks are the same and few recommendations apply to all networks. This topic offers suggestions on how you can use the FortiGate unit to help secure your network against content threats.

IPS signatures

IPS signatures can detect malicious network traffic. For example, the Code Red worm attacked a vulnerability in the Microsoft IIS web server. Your FortiGate's IPS system can detect traffic attempting to exploit this vulnerability. IPS may also detect when infected systems communicate with servers to receive instructions.

IPS recommendations

- Enable IPS scanning at the network edge for all services.
- Use FortiClient endpoint IPS scanning for protection against threats that get into your network.
- Subscribe to FortiGuard IPS Updates and configure your FortiGate unit to receive push updates. This will ensure you receive new IPS signatures as soon as they are available.
- Your FortiGate unit includes IPS signatures written to protect specific software titles from DoS attacks. Enable the signatures for the software you have installed and set the signature action to **Block**.
- You can view these signatures by going to **Security Profiles > Intrusion Prevention** and selecting the **[View IPS Signatures]** link in the right-hand corner of the window.
- Because it is critical to guard against attacks on services that you make available to the public, configure IPS signatures to block matching signatures. For example, if you have a web server, configure the action of web server signatures to **Block**.

Suspicious traffic attributes

Network traffic itself can be used as an attack vector or a means to probe a network before an attack. For example, SYN and FIN flags should never appear together in the same TCP packet. The SYN flag is used to initiate a TCP session while the FIN flag indicates the end of data transmission at the end of a TCP session.

The FortiGate unit has IPS signatures that recognize abnormal and suspicious traffic attributes. The SYN/FIN combination is one of the suspicious flag combinations detected in TCP traffic by the `TCP.BAD.FLAGS` signature.

The signatures that are created specifically to examine traffic options and settings, begin with the name of the traffic type they are associated with. For example, signatures created to examine TCP traffic have signature names starting with TCP.

Application control

While applications can often be blocked by the ports they use, application control allows convenient management of all supported applications, including those that do not use set ports.

Application control recommendations

- Some applications behave in an unusual manner in regards to application control. For more information, see [Application considerations on page 2490](#).
- By default, application control allows the applications not specified in the application control list. For high security networks, you may want to change this behavior so that only the explicitly allowed applications are permitted.

SSL/SSH inspection

Regular web filtering can be circumvented by using `https://` instead of `http://`. By enabling this feature, the FortiGate can filter traffic that is using the HTTPS protocol. This sort of analysis is some times referred to as deep scanning.

Deep Inspection works along the following lines: If your FortiGate unit has the correct chipset it will be able to scan SSL encrypted traffic in the same way that regular traffic can be scanned. The FortiGate firewall will essentially receive the traffic on behalf of the client and open up the encrypted traffic. Once it is finished it re-encrypts the traffic and sends it on to its intended recipient. It is very similar to a man-in-the-middle attack. By enabling this feature, it allows the FortiGate firewall to filter on traffic that is using the SSL encrypted protocol.

The encrypted protocols that can be inspected are:

- HTTPS
- SMTPS
- POP3S
- IMAPS
- FTPS

Before the invention of SSL inspection, scanning regular web traffic can be circumvented by using the prefix `https://` instead of `http://` in the URL. SSL inspection prevents this circumvention. However, because when the encrypted traffic is decrypted it has to be re-encrypted with the FortiGate's certificate rather than the original certificate it can cause errors because the name on the certificate does not match the name on the web site.

At one point deep inspection was something that was either turned on or off. Now individual deep inspection profiles can be created depending on the requirements of the policy. Depending on the Inspection Profile, you can:

- Configure which CA certificate will be used to decrypt the SSL encrypted traffic.
- Configure which SSL protocols will be inspected.
- Configure which ports will be associated with which SSL protocols for the purpose of inspection.
- Configure which websites will be exempt from SSL inspection
- Configure whether or not to allow invalid SSL certificates.
- Configure whether or not SSH traffic will be inspected.

Web rating overrides

This feature allows you to override the FortiGuard Web Filtering. This option allows users to change the rating for a website and control access to the site without affecting the rest of the sites in the original category. More information can be found in [Overriding FortiGuard website categorization](#).

Web profile overrides

This feature allows administrators to grant temporary access to sites that are otherwise blocked by a web filter profile. The temporary access can be granted to a user, user group, or source IP address. The time limit can be set in days, hours, or minutes. See the section on [Web Profile Overrides](#) for more information.

Content inspection and filtering

When the FortiGate unit buffers the packets containing files, email messages, web pages, and other similar files for reassembly before examining them, it is performing content inspection. Traffic inspection, on the other hand, is accomplished by the FortiGate unit examining individual packets of network traffic as they are received.

No two networks are the same and few recommendations apply to all networks. This topic offers suggestions on how you can use the FortiGate unit to help secure your network against threats to content. Be sure to understand the effects of the changes before using the suggestions.

AntiVirus

The FortiGate antivirus scanner can detect viruses and other malicious payloads used to infect machines. The FortiGate unit performs deep content inspection. To prevent attempts to disguise viruses, the antivirus scanner will reassemble fragmented files and uncompress content that has been compressed. Patented Compact Pattern

Recognition Language (CPRL) allows further inspection for common patterns, increasing detection rates of virus variations in the future.

AntiVirus recommendations

- Enable antivirus scanning at the network edge for all services.
- Use FortiClient endpoint antivirus scanning for protection against threats that get into your network.
- Subscribe to FortiGuard AntiVirus Updates and configure your FortiGate unit to receive push updates. This will ensure that new antivirus signatures are loaded onto your FortiGate as soon as they are available.
- Enable the Extended Virus Database if your FortiGate unit supports it.
- Examine antivirus logs periodically. Take particular notice of repeated detections. For example, repeated virus detection in SMTP traffic could indicate a system on your network is infected and is attempting to contact other systems to spread the infection using a mass mailer.
- To conserve system resources, avoid scanning email messages twice. Scan messages as they enter and leave your network or when clients send and retrieve them, rather than both.
- Enable **Treat Windows Executables in Email Attachments as Viruses** if you are concerned about incoming '.exe' files.

FortiGuard web filtering

The web is the most popular part of the Internet and, as a consequence, virtually every computer connected to the Internet is able to communicate using port 80, HTTP. Botnet communications take advantage of this open port and use it to communicate with infected computers. FortiGuard Web Filtering can help stop infections from malware sites and help prevent communication if an infection occurs.

FortiGuard web filtering recommendations

- Enable FortiGuard Web Filtering at the network edge.
- Install the FortiClient application and use FortiGuard Web Filtering on any systems that bypass your FortiGate unit.
- Block categories such as Pornography, Malware, Spyware, and Phishing. These categories are more likely to be dangerous.
- In the Anti-Spam profile, enable **Spam Detection and Filtering** and then enable **IP Address Check**. Many IP addresses used in spam messages lead to malicious sites; checking them will protect your users and your network.

DNS filter

DNS-based web filtering

This feature is similar to the FortiGuard DNS web filtering available in FortiOS 5.2. You can configure DNS web filtering to allow, block, or monitor access to web content according to FortiGuard categories. When DNS web filtering is enabled, your FortiGate must use the FortiGuard DNS service for DNS lookups. DNS lookup requests sent to the FortiGuard DNS service return with an IP address and a domain rating that includes the FortiGuard category of the web page.

If that FortiGuard category is set to block, the result of the DNS lookup is not returned to the requester. If the category is set to redirect, then the address returned to the requester points at a FortiGuard redirect page.

You can also allow access or monitor access based on FortiGuard category.

The following filtering options can be configured in a DNS Filter security profile:

Blocking DNS requests to known Botnet C&C addresses

A new FortiGuard database contains a list of known Botnet C&C addresses. This database is updated dynamically and stored on the FortiGate. This database is covered by FortiGuard web filter licensing; you must have an active FortiGuard web filtering license to use this feature. You can view the botnet lists by going to **System > FortiGuard > Botnet IPs** and **System > FortiGuard > Botnet Domains**.

When you block DNS requests to known Botnet C&C addresses, using IPS, DNS lookups are checked against the Botnet C&C database. All matching DNS lookups are blocked. Matching uses a reverse prefix match, so all sub-domains are also blocked.

To enable blocking of DNS requests to known Botnet C&C addresses, go to **Security Profiles > DNS Filter**, and enable **Block DNS requests to known botnet C&C**. When you do this in FortiOS 5.4.1, you can open a definitions window by clicking on "botnet package."

Static URL filter

The DNS static URL filter allows you to block, exempt, or monitor DNS requests by using IPS to look inside DNS packets and match the domain being looked up with the domains on the static URL filter list. If there is a match the DNS request can be blocked, exempted, monitored, or allowed.

If blocked, the DNS request is blocked and so the user cannot look up the address and connect to the site.

If exempted, access to the site is allowed even if another method is used to block it.

Anti-Spam

Spam is a common means by which attacks are delivered. Users often open email attachments they should not, and infect their own machine. The FortiGate email filter can detect harmful spam and mark it, alerting the user to the potential danger.

Anti-Spam filter recommendations

- Subscribe to the FortiGuard Anti-Spam Filtering service.
- Enable email filtering at the network edge for all types of email traffic.
- Use FortiClient endpoint scanning for protection against threats that get into your network.

Data Leak Prevention

Most security features on the FortiGate unit are designed to keep unwanted traffic out of your network while Data Leak Prevention (DLP) can help you keep sensitive information from leaving your network. For example, credit card numbers and social security numbers can be detected by DLP sensors.

DLP recommendations

- Rules related to HTTP posts can be created, but if the requirement is to block all HTTP posts, a better solution is to use application control or the **HTTP POST Action** option in the web filter profile.
- While DLP can detect sensitive data, it is more efficient to block unnecessary communication channels than to use DLP to examine it. If you don't use instant messaging or peer-to-peer communication in your organization, for example, use application control to block them entirely.

Security profile components

Below is a brief description of the security profiles and their features.



Security Profiles can be configured Globally across multiple VDOMs. See "[Global security profiles across Virtual domains \(VDOMs\)](#)" on page 2587 for more information.

AntiVirus

Your FortiGate unit stores a virus signature database that can identify more than 15,000 individual viruses. FortiGate models that support additional virus databases are able to identify hundreds of thousands of viruses. With a FortiGuard AntiVirus subscription, the signature databases are updated whenever a new threat is discovered.

AntiVirus also includes file filtering. When you specify files by type or by file name, the FortiGate unit will block the matching files from reaching your users.

FortiGate units with a hard drive or configured to use a FortiAnalyzer unit can store infected and blocked files for that you can examine later.

Web filter

Web filtering includes a number of features you can use to protect or limit your users' activity on the web.

FortiGuard Web Filtering is a subscription service that allows you to limit access to web sites. More than 60 million web sites and two billion web pages are rated by category. You can choose to allow or block each of the 77 categories.

URL filtering can block your network users from access to URLs that you specify.

Web content filtering can restrict access to web pages based on words and phrases appearing on the web page itself. You can build lists of words and phrases, each with a score. When a web content list is selected in a web filter profile, you can specify a threshold. If a user attempts to load a web page and the score of the words on the page exceeds the threshold, the web page is blocked.

DNS filter

The FortiGate will inspect DNS traffic to any DNS server, so long as the policy has DNS inspection enabled. The FortiGate will intercept DNS requests, regardless of the destination IP, and redirect it to the FortiGuard Secure DNS server -- this is separate from the FortiGuard DNS server.

The Secure DNS server will resolve and rate the FQDN and send a DNS response which includes both IP and rating of the FQDN back to the FortiGate, where it will handle the DNS response according to the DNS filter profile.

Application control

Although you can block the use of some applications by blocking the ports they use for communications, many applications do not use standard ports to communicate. Application control can detect the network traffic of more than 1,000 applications, improving your control over application communication.

Cloud Access Security Inspection (CASI)

This feature introduces a new security profile called Cloud Access Security Inspection (CASI) that provides support for fine-grained control on popular cloud applications, such as YouTube, Dropbox, Baidu, and Amazon. The CASI profile is applied to a policy much like any other security profile.



Unfortunately CASI does not work when using Proxy-based profiles for AV or Web filtering for example. Make sure to only use Flow-based profiles in combination with CASI on a specific policy.

Intrusion protection

The FortiGate Intrusion Protection System (IPS) protects your network against hacking and other attempts to exploit vulnerabilities of your systems. More than 3,000 signatures are able to detect exploits against various operating systems, host types, protocols, and applications. These exploits can be stopped before they reach your internal network.

You can also write custom signatures tailored to your network.

Anti-spam

FortiGuard Anti-Spam is a subscription service that includes an IP address black list, a URL black list, and an email checksum database. These resources are updated whenever new spam messages are received, so you do not need to maintain any lists or databases to ensure accurate spam detection.

You can use your own IP address lists and email address lists to allow or deny addresses, based on your own needs and circumstances.

Data Leak Prevention

Data Leak Prevention (DLP) allows you to define the format of sensitive data. The FortiGate unit can then monitor network traffic and stop sensitive information from leaving your network. Rules for U.S. social security numbers, Canadian social insurance numbers, as well as Visa, Mastercard, and American Express card numbers are included.

VoIP

The Session Initiation Protocol (SIP) is an IETF application layer signaling protocol used for establishing, conducting, and terminating multi-user multimedia sessions over TCP/IP networks using any media. SIP is often used for Voice over IP (VoIP) calls but can be used for establishing streaming communication between end points.

For more information, see [VoIP Solutions: SIP](#).

ICAP

This module allows for the offloading of certain processes to a separate server so that your FortiGate firewall can optimize its resources and maintain the best level of performance possible.

FortiClient profiles

FortiClient is an all-in-one comprehensive endpoint security solution that extends the power of Fortinet's Advanced Threat Protection (ATP) to end user devices. As the endpoint is the ultimate destination for malware that is seeking credentials, network access, and sensitive information, ensuring that your endpoint security combines strong prevention with detection and mitigation is critical.

The FortiGate provides network security by defining compliance rules for FortiClient endpoints.

For more information, see the [FortiClient 5.4.1 Administration Guide](#).

Proxy options

Proxy Options includes features you can configure for when your FortiGate is operating in proxy mode, including protocol port mapping, block oversized files/emails, and other web and email options.

SSL/SSH inspection

SSL/SSH Inspection (otherwise known as Deep Inspection) is used to scan HTTPS traffic in the same way that HTTP traffic can be scanned. This allows the FortiGate to receive and open up the encrypted traffic on behalf of the client, then the traffic is re-encrypted and sent on to its intended destination.

Individual Deep Inspection profiles can be created, depending on the requirements of the policy. Depending on the profile, you can:

- Configure which CA certificate will be used to decrypt the SSL encrypted traffic
- Configure which SSL protocols will be inspected
- Configure which ports will be associated with which SSL protocols for inspection
- Configure whether or not to allow invalid SSL certificates
- Configure whether or not SSH traffic will be inspected

Security profiles/lists/sensors

A profile is a group of settings that you can apply to one or more firewall policies. Each Security Profile feature is enabled and configured in a profile, list, or sensor. These are then selected in a security policy and the settings apply to all traffic matching the policy. For example, if you create an antivirus profile that enables antivirus scanning of HTTP traffic, and select the antivirus profile in the security policy that allows your users to access the World Wide Web, all of their web browsing traffic will be scanned for viruses.



While you can apply more than one security profile to a firewall policy, it is not recommended that you use flow-based profiles and proxy-based profiles in the same firewall policy.

Because you can use profiles in more than one security policy, you can configure one profile for the traffic types handled by a set of firewall policies requiring identical protection levels and types, rather than repeatedly configuring those same profile settings for each individual security policy.

For example, while traffic between trusted and untrusted networks might need strict protection, traffic between trusted internal addresses might need moderate protection. To provide the different levels of protection, you might configure two separate sets of profiles: one for traffic between trusted networks, and one for traffic between trusted and untrusted networks.

Inspection modes

You can select one of two inspection modes from the **System > Settings** page to control the security profile inspection mode for your FortiGate or VDOM.

- **Proxy-based inspection**, that reconstructs content passing through the FortiGate unit and inspects the content for security threats, or
- **Flow-based inspection**, that takes a snapshot of content packets and uses pattern matching to identify security threats in the content.

Each inspection component plays a role in the processing of traffic en route to its destination. Having control over flow and proxy mode is helpful if you want to be sure that only flow inspection mode is used (and that proxy inspection mode is not used). In most cases proxy mode is preferred because more security profile features are available and more configuration options for these individual features are available. Yet, some implementations may require all security profile scanning to only use flow mode. In this case, you can set your FortiGate to flow mode knowing that proxy mode inspection will not be used. While both modes offer significant security, proxy-based provides more features and flow-based is designed to optimize performance.

This section addresses the following topics:

Proxy-based inspection

If a FortiGate or VDOM is configured for proxy-based inspection, then a mixture of flow-based and proxy-based inspection occurs. Traffic initially encounters the IPS engine, which applies single-pass IPS, Application Control, and CASI, if configured in the firewall policy accepting the traffic.

The traffic is then sent for proxy-based inspection. Proxy-based inspection extracts and caches content, such as files and web pages, from a content session and inspects the cached content for threats. Content inspection takes place in the following order: **VoIP inspection, DLP, AntiSpam, Web Filtering, AntiVirus, and ICAP**.

If no threat is found, the proxy relays the content to its destination. If a threat is found, the proxy can block the threat and send a replacement message in its stead. The proxy can also block VoIP traffic that contains threats.

Transparent web proxy mode

In proxy mode, FortiOS 5.6 functions just like FortiOS 5.4 with the addition of the new Transparent Web Proxy mode. See New Operating mode for Transparent web proxy in [What's New in FortiOS 5.6](#).

Flow-based inspection

Flow-based inspection identifies and blocks security threats in real time as they are identified using single-pass Direct Filter Approach (DFA) pattern matching to identify possible attacks or threats.

If a FortiGate or a VDOM is configured for flow-based inspection, depending on the options selected in the firewall policy that accepted the session, flow-based inspection can apply **IPS, Application Control, Web Filtering, DLP, and AntiVirus**. Flow-based inspection is all done by the IPS engine and, as you would expect, no proxying is involved.

All of the applicable flow-based security modules are applied simultaneously in one single pass, and pattern matching is offloaded and accelerated by CP8 or CP9 processors. **IPS, Application Control**, flow-based **Web**

Filtering, and flow-based **DLP** filtering happen together. Flow-based **AntiVirus** scanning caches files during protocol decoding and submits cached files for virus scanning while the other matching is carried out.

Flow-based inspection typically requires fewer processing resources than proxy-based inspection and does not change packets, unless a threat is found and packets are blocked. Flow-based inspection cannot apply as many features as proxy inspection. For example, flow-based inspection does not support client comforting and some aspects of replacement messages.

In FortiOS 5.6, flow-based inspection requires the new **NGFW mode**.

Changing between proxy and flow mode

You can see which inspection mode your FortiGate is using by looking at the **System Information** widget on your **Dashboard**.

To change inspection modes, go to **System > Settings** and scroll down to **Inspection Mode**. You can select Flow-based to operate in Flow mode or Proxy to operate in Proxy mode.

When you select **Flow-based**, all proxy mode profiles are converted to flow mode, removing any proxy settings. As well proxy mode only features (for example, Web Application Profile) are removed from the GUI.

In addition, selecting **Flow-based** inspection will cause the **Explicit Web Proxy** and **Explicit FTP Proxy** features to be removed from the GUI and the CLI. This includes Explicit Proxy firewall policies.

When you select **Flow-based** you can only configure Virtual Servers (under **Policy & Objects > Virtual Servers**) with Type set to HTTP, TCP, UDP, or IP.

If required, you can change back to proxy mode through the **System > Settings** page.

If your FortiGate has multiple VDOMs, you can set the inspection mode independently for each VDOM. Use the top left drop-down menu to go to **Global > System > VDOM**. Click **Edit** for the VDOM you wish to change and select the **Inspection Mode**.



Switching to flow-based inspection also turns off WAN Optimization, Web Caching, the Explicit Web Proxy, and the Explicit FTP Proxy making sure that no proxying can occur.

From the GUI, you can only configure antivirus and web filter security profiles in proxy mode. From the CLI you can configure flow-based antivirus profiles, web filter profiles and DLP profiles and they will appear on the GUI and include their inspection mode setting. Also, flow-based profiles created when in flow mode are still available when you switch to proxy mode.

NGFW profile-based and NGFW policy-based modes

When you select **Flow-based** as the **Inspection Mode**, you have the option in FortiOS 5.6 to select an **NGFW Mode**. **NGFW Profile-based** mode works the same as flow-based mode did in FortiOS 5.4

When selecting NGFW policy-based mode you can also select the SSL/SSH Inspection mode that is applied to all policies.

In the new **NGFW Policy-based** mode, you add applications and web filtering profiles directly to a policy without having to first create and configure Application Control or Web Filtering profiles. See [NGFW Policy Mode on page 1](#).

When you change to flow-based inspection, all proxy mode profiles are converted to flow mode, removing any proxy settings. And proxy-mode only features (for example, Web Application Profile) are removed from the GUI.

If your FortiGate has multiple VDOMs, you can set the inspection mode independently for each VDOM. Go to **System > VDOM**. Click **Edit** for the VDOM you wish to change and select the **Inspection Mode**.

CLI syntax

The following CLI commands can be used to configure inspection and NGFW (called "policy" in the CLI) modes:

```
config system settings
  set inspection-mode {proxy | flow}
  set policy-mode {standard | ngfw}
end
```

Comparison of inspection types

The tables in this section show how different security features map to different inspection types and present the strengths and weaknesses of proxy- vs. flow-based inspection.

Security profile features mapped to inspection mode

The table below lists FortiOS security profile features and shows whether they are available in flow-based or proxy-based inspection modes.

Security Profile Feature	Flow-based inspection	Proxy-based inspection
AntiVirus	x	x
Web Filter	x	x
DNS Filter	x	x
Application Control	x	x
Intrusion Protection	x	x
Anti-Spam		x
Data Leak Protection		x
VoIP		x
ICAP		x
Web Application Firewall		x

Security Profile Feature	Flow-based inspection	Proxy-based inspection
FortiClient Profiles	x	x
Proxy Options	x	x
SSL Inspection	x	x
SSH Inspection		x
Web Rating Overrides	x	x
Web Profile Overrides		x

Individual security profile considerations

In flow mode, AntiVirus and Web Filter profiles only include flow-mode features. Web filtering and virus scanning are still done with the same engines and to the same accuracy, but some inspection options are limited or not available in flow mode. Application control, intrusion protection, and FortiClient profiles are not affected when switching between flow and proxy mode.

Application control uses flow-based inspection; if you apply an additional security profile to your traffic that is proxy-based, the connection will simply timeout rather than display the warning, or replacement, message. However, Application Control will still function.

Even though VoIP profiles are not available from the GUI in flow mode, the FortiGate can process VoIP traffic. In this case the appropriate session helper is used (for example, the SIP session helper).

Setting flow or proxy mode doesn't change the settings available from the CLI. However, when in flow mode you can't save security profiles that are set to proxy mode.

You can also add proxy-only security profiles to firewall policies from the CLI. So, for example, you can add a VoIP profile to a security policy that accepts VoIP traffic. This practice isn't recommended because the setting will not be visible from the GUI.

If you set flow-based to use external servers for FortiWeb and FortiMail you must use the CLI to set a Web Application Firewall profile or Anti-Spam profile to external mode and add the Web Application Firewall profile or AntiSpam profile to a firewall policy.

Proxy mode and flow mode antivirus and web filter profile options

The following tables list the antivirus and web filter profile options available in proxy and flow modes.

Antivirus features in proxy and flow mode

Feature	Proxy	Flow
Scan Mode (Quick or Full)	no	yes
Detect viruses (Block or Monitor)	yes	yes

Feature	Proxy	Flow
Inspected protocols	yes	no (all relevant protocols are inspected)
Inspection Options	yes	yes (not available for quick scan mode)
Treat Windows Executables in Email Attachments as Viruses	yes	yes
Send Files to FortiSandbox Appliance for Inspection	yes	yes
Use FortiSandbox Database	yes	yes
Include Mobile Malware Protection	yes	yes

Web filter features in proxy and flow mode

Feature	Proxy	Flow
FortiGuard category based filter	yes	yes (show, allow, monitor, block)
Category Usage Quota	yes	no
Allow users to override blocked categories (on some models)	yes	no
Search Engines	yes	no
Enforce 'Safe Search' on Google, Yahoo!, Bing, Yandex	yes	no
Restrict YouTube Access	yes	no
Log all search keywords	yes	no
Static URL Filter	yes	yes
Block invalid URLs	yes	no
URL Filter	yes	yes
Block malicious URLs discovered by FortiSandbox	yes	yes
Web Content Filter	yes	yes
Rating Options	yes	yes

Feature		Proxy	Flow
	Allow websites when a rating error occurs	yes	yes
	Rate URLs by domain and IP Address	yes	yes
	Block HTTP redirects by rating	yes	no
	Rate images by URL	yes	no
Proxy Options		yes	no
	Restrict Google account usage to specific domains	yes	no
	Provide details for blocked HTTP 4xx and 5xx errors	yes	no
	HTTP POST Action	yes	no
	Remove Java Applets	yes	no
	Remove ActiveX	yes	no
	Remove Cookies	yes	no
	Filter Per-User Black/White List	yes	no

AntiVirus scanning differences between versions of FortiOS 5.x

In FortiOS 5.0, 5.2, 5.4, 5.6 and 6.0, there are several AntiVirus (AV) scanning inspection modes available. FortiOS 5.0 includes proxy and flow-based virus scanning. FortiOS 5.2 also uses proxy-based and flow-based scanning, but the flow-based mode in FortiOS 5.2 uses a new approach to flow-based scanning (that is sometimes called deepflow or deep flow scanning). FortiOS 5.4 and onward offer another flow-based mode, quick mode, to inspect traffic efficiently.

The databases used for AV scanning does not change from proxy to flow mode unless quick mode is enabled. In flow-based quick mode, a compact antivirus database is used.

AntiVirus scanning examines files in HTTP, HTTPS, email, and FTP traffic for threats as they pass through your FortiGate. If the traffic contains compressed files, they are also examined. Go to the SysAdmin Note on the Fortinet Cookbook site for detailed information on [supported compression formats](#) in antivirus scanning.

If the AV scanner finds a threat such as a virus or some other malware, FortiOS protects your network by blocking the file.

FortiOS includes a number of AntiVirus features that make virus scanning more user-friendly. One of these features, called replacement messages, sends a customizable message to anyone whose file is blocked by AV scanning, to explain what happened and why. Other features make communication between the client and the server more seamless. The availability of these changes depending on the inspection mode.

Proxy-based AV scanning

Proxy-based AV scanning is the most feature-rich AV scanning mode. This mode uses a proxy to manage the communication between client and server. The proxy extracts content packets from the data stream as they arrive and buffers the content until the complete file is assembled. Once the file is whole, the AV scanner examines the file for threats. If no threats are found, the file is sent to its destination. If a threat is found, the file is blocked.

Because proxy-based scanning is applied to complete files, including compressed files, it provides very effective threat detection. Proxy-based scanning also supports a full range of features, including replacement messages and client comforting, making proxy-based scanning the most user friendly inspection mode. In addition the proxy manages the communication between the client and the server, improving the user experience. For example, in flow mode if a virus is found, the last part of the file is not downloaded and the connection just times out and the user cannot tell what is going on. In proxy mode, the users gets a message about the file being blocked.

Proxy-based scanning inspects all files under the oversized threshold. Since the FortiGate unit has a limited amount of memory, files larger than a certain size do not fit within the memory buffer. The default buffer size is 10 MB. You can use the `uncompssize limit` CLI command to adjust the size of this memory buffer. Files larger than the threshold are passed to the destination without scanning. You can use the **Oversized File/Email** setting in **Security Profiles > Proxy Options** to block files larger than the antivirus buffer if allowing files that are too large to be scanned is an unacceptable security risk.

During the buffering and scanning procedure, the client must wait. With a default configuration, the file is released to the client only after it is scanned. You can enable client comforting in the **Proxy Options** security profile to feed the client a trickle of data to prevent them from possibly thinking the transfer is stalled and consequently canceling the download.

Flow-based AV scanning

Although the name "flow-based scanning" is used in FortiOS 5.0, 5.2, 5.4, and 5.6, the different versions handle this mode in very different ways.

Flow AV in FortiOS 5.4 and 5.6

In FortiOS 5.4 and 5.6, there are two modes available for flow-based virus scanning: **Quick** and **Full** scan mode. Full mode is the same as flow-based scanning in FortiOS 5.2 (see below). Quick mode uses a compact antivirus database and advanced techniques to improve performance. You can designate quick or full scan mode when configuring the antivirus profile in the GUI. Alternatively, use the following CLI command to enable quick or full mode:

```
config antivirus profile
edit <profile>
set scan-mode {quick | full}
end
```

Flow AV in FortiOS 5.2 (deepflow or deep flow)

FortiOS 5.2 introduced a new type of flow-based AV scanning, that is sometimes called deepflow or deep flow, and that takes a hybrid approach where content packets are buffered while simultaneously being sent to their destination. When all of the files packets have been collected and buffered, but before the final packet is delivered, the buffered file is scanned. If a threat is found, the last packet is blocked and the client application has to deal with not getting the completed file. If no threat is found the final packet is sent and the user gets their file.

Deepflow AV scanning is as good as proxy-based AV scanning at detecting threats. There may be a small performance advantage over proxy-based AV as files get larger based on the difference between sending the

whole file after analysis and just sending the last packet. Deepflow's most notable limitation is that, just like the flow-based AV in 5.0, it does not support many of the user-friendly features provided by proxy-based AV.

Flow AV in FortiOS 5.0

In FortiOS 5.0, flow-based AV scanning examines the content of individual data packets as they pass through the FortiGate. There is no proxy involved so packets are not changed by the proxy and files are not buffered for analysis. Potentially less memory and CPU resources are used, resulting in a potential performance increase compared to using proxy-based mode. FortiOS 5.0 flow-based AV scanning is also not limited by file size.

Flow AV uses the IPS engine and the AV database and is effective at many kinds of threat detection; however, because it can only analyze what is in an individual packet rather than a complete file, flow-based scanning cannot detect some types of malware, including polymorphic code. Malware in documents, compressed files, and some archives are also less likely to be detected.

Flow AV does not actually block files, it stops delivering a file's packets once a threat has been detected. This means that parts of the file may already have been delivered when the threat has been detected and the recipient application is responsible for dealing with the partially complete content.

In addition flow AV can be less user friendly. Replacement messages are not supported and clients may have to wait for sessions to time out without knowing why content has been blocked.

AntiVirus

This section describes how to configure the antivirus options. From an antivirus profile you can configure the FortiGate unit to apply antivirus protection to HTTP, FTP, IMAP, POP3, SMTP, and NNTP sessions. If your FortiGate unit supports SSL/SSH content scanning and inspection, you can also configure antivirus protection for HTTPS, IMAPS, POP3S, SMTPS, and FTPS sessions.

In many cases you can just customize the default antivirus profile and apply it to the security policy that accepts the traffic to be virus scanned. You can also create custom antivirus profiles if want to apply different types of virus protection to different traffic.

This Handbook chapter includes [Inside FortiOS: AntiVirus](#) providing readers an overview of the features and benefits of key FortiOS components.

For readers needing to delve into greater detail, we provide the following topics:

- [Antivirus concepts](#)
- [Enabling AntiVirus scanning](#)
- [Testing your antivirus configuration](#)
- [Example Scenarios](#)

Antivirus concepts

The word “antivirus” refers to a group of features that are designed to prevent unwanted and potentially malicious files from entering your network. These features all work in different ways, which include checking for a file size, name, or type, or for the presence of a virus or grayware signature.

The antivirus scanning routines your FortiGate unit uses are designed to share access to the network traffic. This way, each individual feature does not have to examine the network traffic as a separate operation, and the overhead is reduced significantly. For example, if you enable file filtering and virus scanning, the resources used to complete these tasks are only slightly greater than enabling virus scanning alone. Two features do not require twice the resources.

Antivirus scanning examines files for viruses, worms, trojans, and other malware. The antivirus scan engine has a database of virus signatures it uses to identify infections. If the scanner finds a signature in a file, it determines that the file is infected and takes the appropriate action.

Malware threats

Malware is the general term covering all the different types of threats to your computer safety such as:

- [Viruses](#)
- [Worms](#)
- [Trojan horses](#)
- [Ransomware](#)
- [Scareware](#)
- [Spyware](#)
- [Adware](#)
- [Botnets](#)

- [Phishing](#)
- [Grayware](#)

Viruses

Viruses are self-replicating code that install copies of themselves into other programs or data files for boot sectors of storage devices. Viruses can often carry a “payload” which performs some undesirable function. These functions can include but are not limited to:

- Stealing drive space
- Stealing CPU cycles
- Accessing private information
- Corrupting data
- Digital defacement or vandalism
- Spamming contact lists

Worms

A worm is a piece of standalone computer code that replicates itself in order to spread to other computers. It normally uses a computer network to spread itself, using security vulnerabilities on the target computer or network to propagate. Unlike a virus, it does not attach itself to an existing file. Even if there is no payload, worms consume resources such as bandwidth and storage space just through their act of replication.

Trojan horses

A Trojan horse, or Trojan is malware that is defined by its delivery method. Through the use of social engineering, or some other method, the code is installed on a system by a valid user of the system and like the original Trojan horse there is something more than advertised within the software. Trojans, unlike worms or viruses are generally non-self-replicating. The most common payload of a Trojan is the setting up of a “backdoor” control mechanism to the system that it is installed on.

Ransomware

Ransomware is a type of malware that, as the name implies, hold the system ransom until payment of some kind is made. It does this by restricting access to the legitimate owner of the system either by encrypting files or locking the system. Usually, a message of some kind is displayed with the demands. Upon payment a utility or key is sent to the user to unlock the system.

Scareware

Scareware comes in two main flavours; the first tries to convince the user that his computer is infected with some non-existent malware, scaring the user into purchasing the author’s virus removal utility. The utility is nonfunctional or some additional form of malware.

The second form tries to convince the user that the computer has been or is being used for an illegal act, such as being part of a botnet or storing child pornography. Again, the objective is to scare the user into paying to cure something that is not really there.

Spyware

Spyware is used by its authors to collect information about the user and its computer without the user's knowledge. The end result can be as benign as being better able to target ads, to as criminal as key loggers designed to record account ids and passwords of bank accounts and forward them off to the authors.

Adware

Adware is not malware per se. It is merely any software that produces advertisements in order to generate revenue for its author. While a lot of people find this inconvenient or irritating, it is not malware. As such, it is not blocked by the antivirus software for being malware.

Software that has adware built into it will be blocked if it has malware in it.

Botnets

A botnet is a network of Internet connected computers that have been covertly usurped to forward transmissions to other computers on the Internet on behalf of a "master". These transmissions can be minimally damaging, such as spam, or they can critically impact a target as when used to launch a Distributed Denial of Service attack.

Any such computer is referred to as a zombie - in effect, a computer "robot" or "bot" that serves the wishes of some master spam or virus originator. Most computers compromised in this way are home-based.

According to a report from Russian-based Kaspersky Labs, botnets -- not spam, viruses, or worms -- currently pose the biggest threat to the Internet. A report from Symantec came to a similar conclusion.

See also: [Botnet protection](#).

Phishing

Phishing is a social engineering technique that is used to obtain sensitive and confidential information by masquerading as a communication from a trusted entity such as a well-known institution, company, or website. Usually, the malware is not in the communication itself but in the links within the communication.

Grayware

Grayware programs are unsolicited software programs installed on computers, often without the user's consent or knowledge. Grayware programs are generally considered an annoyance, but they can also cause system performance problems or be used for malicious purposes.

AntiVirus scanning order

The antivirus scanning function includes various modules and engines that perform separate tasks.

FortiOS has two different modes of antivirus scanning: **proxy-based** and **flow-based**. The reasons for the different modes are performance and granularity. In just about everything relating to security there is a constant balancing act going on. As the level of security and comprehensiveness increases, there is by necessity a decrease in either convenience or performance or both. The increase in processing to scan for more threats requires more resources; resources that are a finite supply on the hardware. Granularity can sometimes be used to mitigate performance impact by scanning for a smaller subset of traffic but this is only recommended when that smaller subset of traffic is the only traffic going through the firewall.

If the traffic on the device is slight, then the impact on the performance will hardly be noticeable. But if the unit is working close to capacity in terms of traffic and there are a lot of files coming through, then there might be a noticeable decline in the performance.

While both modes offer significant security, proxy-based is weighted towards being more thorough and easily configurable, while flow-based is designed to optimize performance.



See [Antivirus scanning differences in FortiOS 5.0, 5.2, 5.4, and 5.6](#) in the [Inspection Modes](#) section for more details on flow vs. proxy inspection modes on your FortiGate.

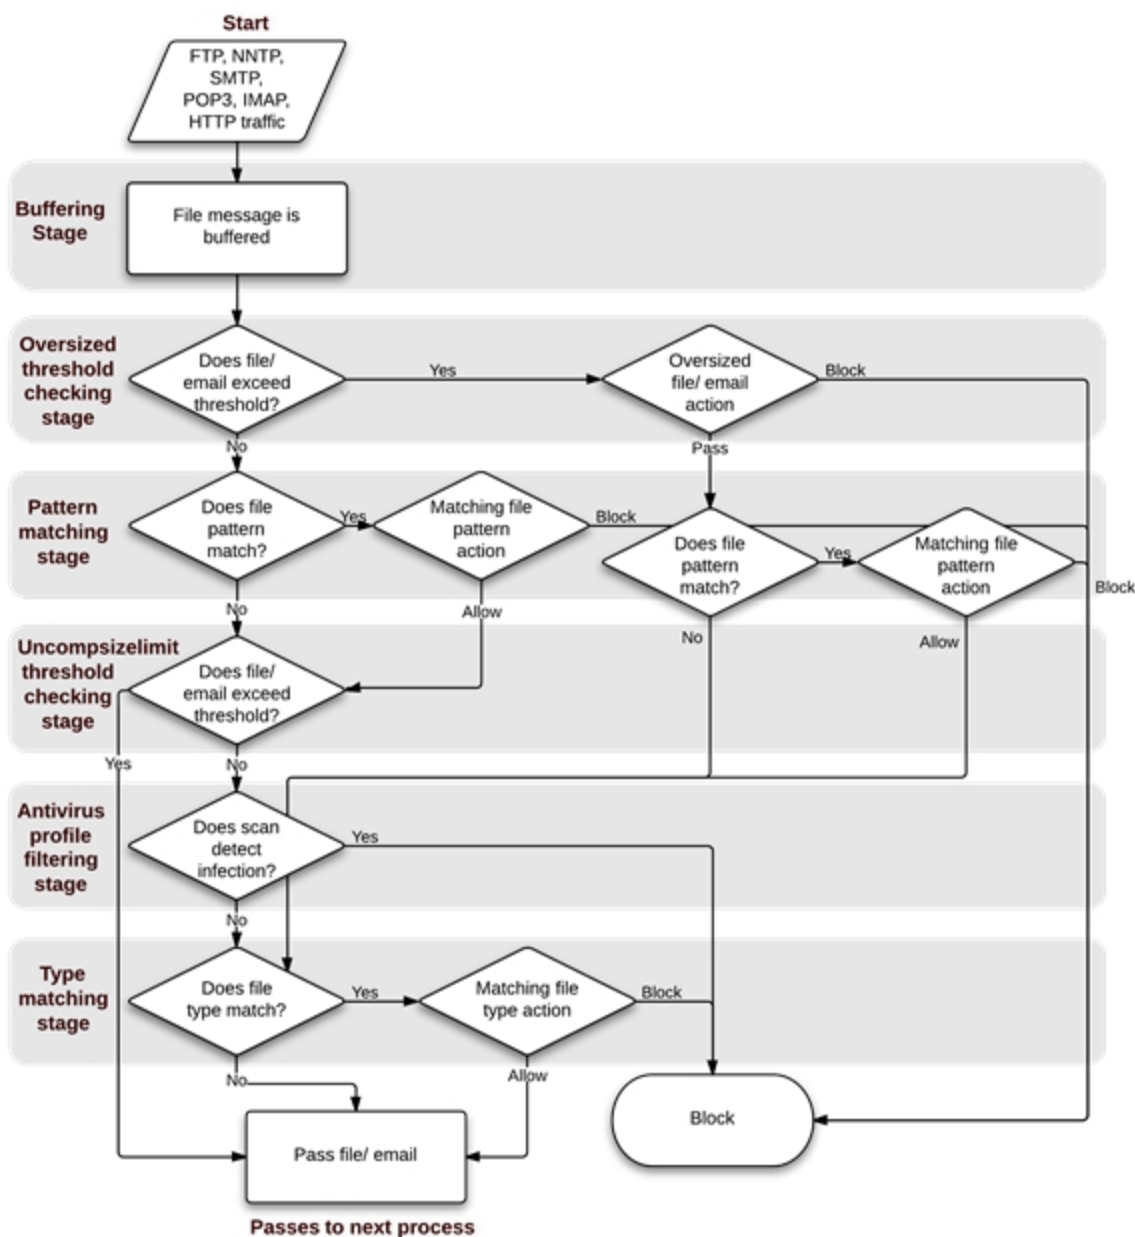
Proxy-based antivirus scanning order

The following figure illustrates the antivirus scanning order when using proxy-based scanning. The first check for oversized files/email is to determine whether the file exceeds the configured size threshold. The `uncompsizelimit` check is to determine if the file can be buffered for file type and antivirus scanning. If the file is too large for the buffer, it is allowed to pass without being scanned. For more information, see the `config antivirus service` command. The antivirus scan includes scanning for viruses, as well as for grayware and heuristics, if enabled.



File filtering includes file pattern and file type scans which are applied at different stages in the antivirus process.

Antivirus scanning order when using the normal, extended, or extreme database

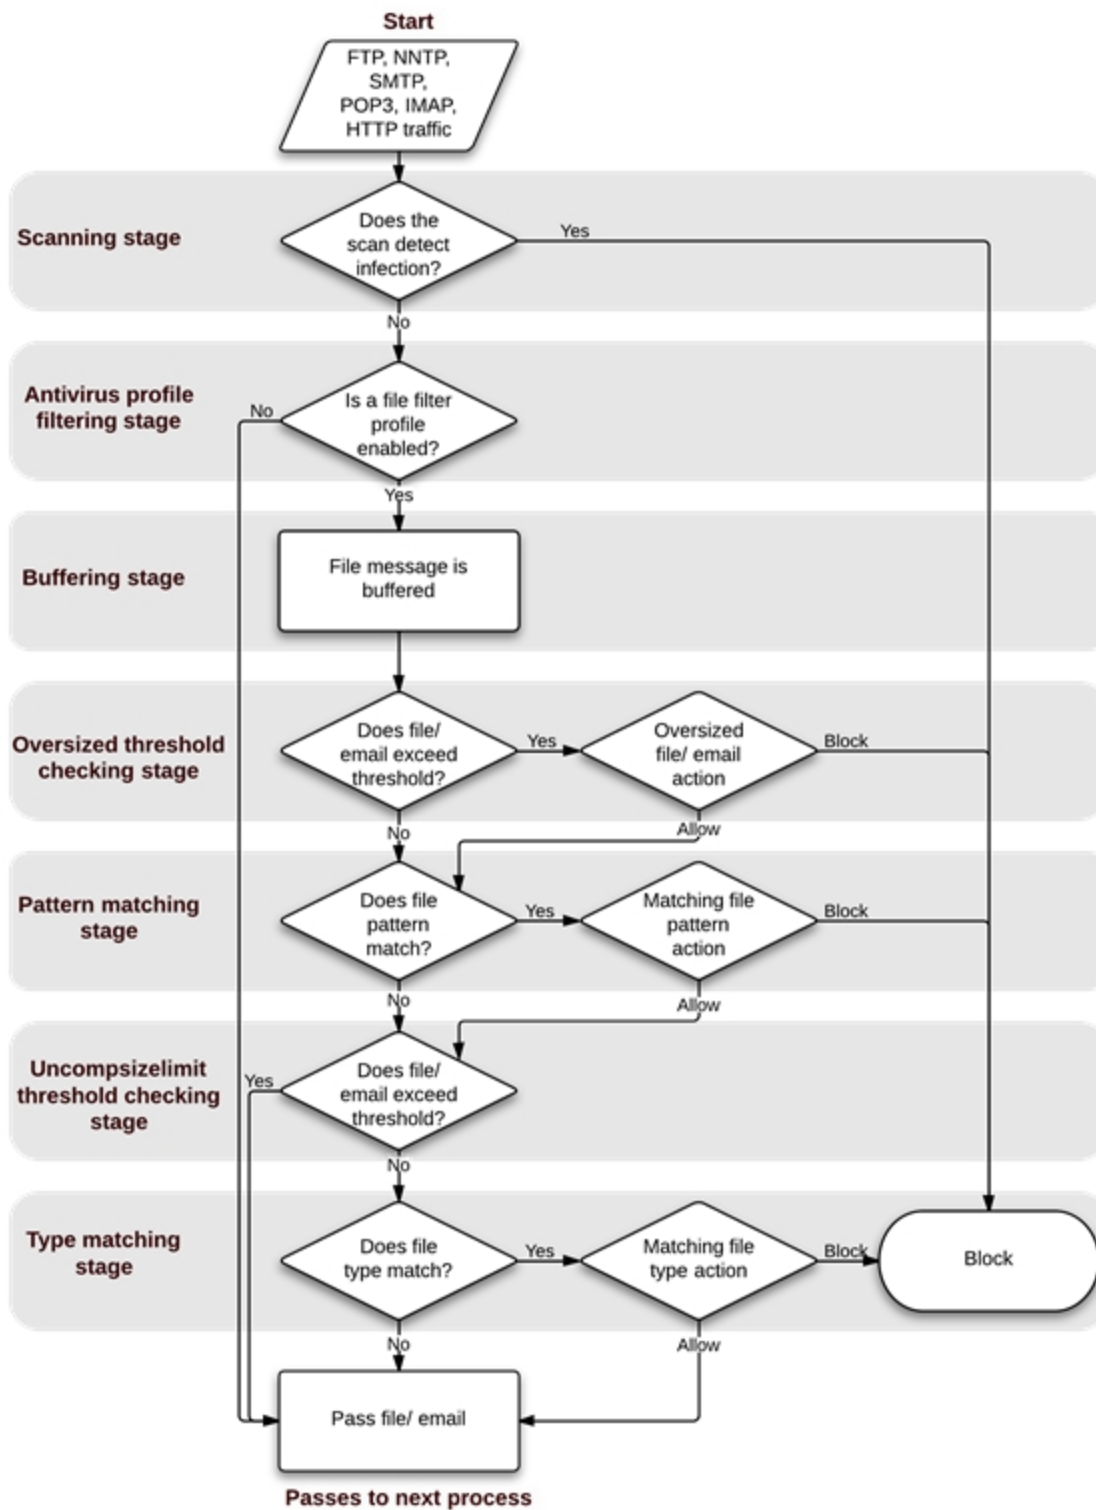


If a file fails any of the tasks of the antivirus scan, no further scans are performed. For example, if the file `fakefile.EXE` is recognized as a blocked file pattern, the FortiGate unit will send the end user a replacement message, and delete or quarantine the file. The unit will not perform virus scan, grayware, heuristics, and file type scans because the previous checks have already determined that the file is a threat and have dealt with it.

Flow-based antivirus scanning order

The following figure illustrates the antivirus scanning order when using flow-based scanning (i.e. the flow-based database). The antivirus scan takes place before any other antivirus-related scan. If file filter is not enabled, the

file is not buffered. The antivirus scan includes scanning for viruses, as well as for grayware and heuristics if they are enabled.



AntiVirus databases

The antivirus scanning engine relies on a database of virus signatures to detail the unique attributes of each infection. The antivirus scan searches for these signatures, and when one is discovered, the FortiGate unit determines the file is infected and takes action.

All FortiGate units have the normal antivirus signature database but some models have additional databases you can select for use. Which you choose depends on your network and security needs.

Normal	Includes viruses currently spreading as determined by the FortiGuard Global Security Research Team. These viruses are the greatest threat. The Normal database is the default selection and it is available on every FortiGate unit.
Extended	Includes the normal database in addition to recent viruses that are no-longer active. These viruses may have been spreading within the last year but have since nearly or completely disappeared.
Extreme	Includes the extended database in addition to a large collection of 'zoo' viruses. These are viruses that have not spread in a long time and are largely dormant today. Some zoo viruses may rely on operating systems and hardware that are no longer widely used.

If your FortiGate unit supports extended, extreme, or flow-based virus database definitions, you can select the virus database most suited to your needs.

If you require the most comprehensive antivirus protection, enable the extended virus database. The additional coverage comes at a cost, however, because the extra processing requires additional resources.

To change the antivirus database

Use the CLI to run the following commands:

```
config antivirus settings
    set default-db extended
end
```

AntiVirus techniques

The first three antivirus features in the list below work in sequence to efficiently scan incoming files and offer your network optimal antivirus protection. The first two features have specific functions, the third, heuristics, protects against new or previously unknown virus threats.

- **Virus scan**

If the file passes the file pattern scan, the FortiGate unit applies a virus scan to it. The virus definitions are kept up-to-date through the FortiGuard Distribution Network (FDN).

- **Grayware protection**

If the file passes the virus scan, it can be checked for grayware. Grayware scanning is an optional function and must be enabled in the CLI if it is to be scanned for along with other malware. Grayware cannot be scanned for on

its own. While done as a separate step, antivirus scanning must be enabled as well.

To enable grayware detection enter the following in the CLI:

```
config antivirus settings
  set grayware enable
end
```

To disable grayware detection enter the following in the CLI:

```
config antivirus settings
  set grayware disable
end
```

Grayware signatures are kept up to date in the same manner as the antivirus definitions.

• Heuristics

After an incoming file has passed the grayware scan, it is subjected to the heuristics scan. The FortiGate heuristic antivirus engine, if enabled, performs tests on the file to detect virus-like behavior or known virus indicators. In this way, heuristic scanning may detect new viruses, but may also produce some false positive results. You configure heuristics from the CLI.

To set heuristics, enter the following in the CLI:

```
config antivirus heuristic
  set mode {pass | block |disable}
end
```

- “block” enables heuristics and any files determined to be malware are blocked from entering the network.
- “pass” enables heuristics but any files determined to be malware are still allowed to pass through to the recipient.
- “disable” turns off heuristics.

• FortiGuard AntiVirus

The FortiGuard Antivirus services are included in the regular FortiGuard subscription and include automatic updates of antivirus engines and definitions as well as a DNS black list (DNSBL) through the FortiGuard Distribution Network (FDN).

Current information about your subscription and version numbers can be found at **System > FortiGuard**. This page will also allow the configuration of connections to the FortiGuard Center and how often to check for updates to the antivirus files.



Updating antivirus definitions can cause a short disruption of traffic being scanned while the FortiGate unit applies the new signature database. Schedule updates for time periods when traffic is light to minimize disruption.

• Botnet protection

A botnet is a network of Internet connected computers that have been covertly usurped to forward transmissions to other computers on the Internet on behalf of a “master”. These transmissions can be minimally damaging, such as spam, or they can critically impact a target as when used to launch a Distributed Denial of Service attack.

Any such computer is referred to as a zombie - in effect, a computer “robot” or “bot” that serves the wishes of some master spam or virus originator. Most computers compromised in this way are home-based.

The latest botnet database is available from FortiGuard. To see the version of the database and display its contents, go to **System > FortiGuard > AntiVirus >** and you will see data for **Botnet IPs** and **Botnet Domains**. You can also block, monitor, or allow outgoing connections to botnet sites for each FortiGate interface.

• Quarantine / Source IP ban

As of FortiOS 5.2, quarantine was a place where traffic content was held in storage where it couldn't interact with the network or system. This was removed, but the term quarantine was kept to describe keeping selected source IPs from interacting with the network and protected systems. This source IP ban is kept in the kernel rather than in any specific application engine and can be queried by APIs. The features that can use the APIs to access and use the banned source IP addresses are antivirus, DLP, DoS and IPS. Both IPv4 and IPv6 version are included in this feature.

You quarantine a source address through the GUI. Go to **FortiView > Sources**. Right-click on the source address you wish to quarantine and select **Quarantine Source Address**. You can set the duration of the quarantine in days, hours, minutes, or seconds. A User Quarantine ban can be removed in **Monitor > User Quarantine Monitor**.

To configure the AntiVirus security profile to add the source IP address of an infected file to the quarantine or list of banned source IP addresses in the CLI:

```
config antivirus profile
  edit <name of profile>
    config nac-quar
      set infected quar-src-ip
      set expiry 5m
    end
```

If the `quar-src-ip` action is used, the additional variable of expiry time will become available. This variable determines for how long the source IP address will be blocked. In the CLI the option is called `expiry` and the duration is in the format `<###d##h##m>`. The maximum days value is 364. The maximum hour value is 23 and the maximum minute value is 59. The default is 5 minutes.

FortiGuard AntiVirus updates

To ensure that your system receives the most protection available, all virus definitions and signatures are updated regularly through the FortiGuard AntiVirus services. To configure this feature, go to **System > FortiGuard**. Under **AntiVirus & IPS Updates**, enable **Scheduled Updates**. From here you can schedule updates to occur on a consistent weekly, daily, or even hourly basis.



Updating antivirus definitions can cause a short disruption of traffic being scanned while the FortiGate unit applies the new signature database. Schedule updates for time periods when traffic is light to minimize disruption.

FortiSandbox

Not every piece of malware has a signature. This is especially true of new malware and variations on existing malware. FortiOS can upload suspicious files to FortiSandbox for sandbox inspection. When a FortiGate uses sandbox inspection, files are sent to the FortiSandbox. Then the FortiSandbox uses virtual machines (VMs)

running different operating systems to test the file, to determine if it is malicious. If the file exhibits risky behavior, or is found to contain a virus, a new signature can be added to both the local FortiGate malware database and the FortiGuard AntiVirus signature database.

A file is deemed suspicious when it does not contain a known threat but has characteristics that suggest it may be malware. The characteristics that determine if a file is suspicious are updated by Fortinet to reflect the current threat climate.

FortiSandbox is available as a physical or virtual appliance (FortiSandbox Appliance), or as a cloud advanced threat protection service integrated with FortiGate (FortiCloud).

To configure an AntiVirus profile to send files to FortiSandbox, first verify that your FortiSandbox appliance is configured or that your FortiCloud account is active. Then go to **Security Profiles > AntiVirus** and enter the desired **Inspection Options**.

Sending files to the FortiSandbox appliance or to FortiSandbox Cloud does not block files immediately. Instead, the files assist in the discovery of new threats and the creation of new signatures to be added to the global FortiGuard AntiVirus database. Files deemed malicious are also immediately added to a custom Malware Package which is downloaded by the FortiGate every two minutes for live detection.

The **Advanced Threat Protection Statistics** dashboard widget displays the number of files that your FortiGate unit has uploaded or submitted to FortiSandbox. To see FortiSandbox statistics for the last 7 days, go to **Fortinet Security Fabric > Settings**.

Option for "Suspicious Files Only" for FortiSandbox submissions

Beginning in FortiOS 6.0.1, FortiGates can use the FortiSandbox Cloud service as part of the AntiVirus subscription. In order to reduce client upload bandwidth usage and general load on the FortiSandbox service, a new "Suspicious Files Only" upload option has been added to the AntiVirus profile, which previously only had "None" and "All Supported Files".

In order to enforce best practices, "None" is now the default.

Syntax

```
config antivirus profile
  edit <profile name>
    set ftgd-analytics [disable|suspicious|everything]
  end
```

Client comforting

When proxy-based antivirus scanning is enabled, the FortiGate unit buffers files as they are downloaded. Once the entire file is captured, the FortiGate unit scans it. If no infection is found, the file is sent along to the client. The client initiates the file transfer and nothing happens until the FortiGate finds the file clean, and releases it. Users can be impatient, and if the file is large or the download slow, they may cancel the download, not realizing that the transfer is in progress.

The client comforting feature solves this problem by allowing a trickle of data to flow to the client so they can see the file is being transferred. The default client comforting transfer rate sends one byte of data to the client every ten seconds. This slow transfer continues while the FortiGate unit buffers the file and scans it. If the file is infection-free, it is released and the client will receive the remainder of the transfer at full speed. If the file is infected, the FortiGate unit caches the URL and drops the connection. The client does not receive any notification of what happened because the download to the client had already started. Instead, the download stops and the user is left with a partially downloaded file.

If the user tries to download the same file again within a short period of time, the cached URL is matched and the download is blocked. The client receives the Infection cache message replacement message as a notification that the download has been blocked. The number of URLs in the cache is limited by the size of the cache.



Client comforting can send unscanned and potentially infected content to the client. You should only enable client comforting if you are prepared to accept this risk. Keeping the client comforting interval high and the amount low will reduce the amount of potentially infected data that is downloaded.

Client comforting is available for HTTP and FTP traffic. If your FortiGate unit supports SSL content scanning and inspection, you can also configure client comforting for HTTPS and FTPS traffic.

Enable and configure client comforting

1. Go to **Security Profiles > Proxy Options**.
2. Select a Proxy Options profile and choose **Edit**, or select **Create New** to make a new one.
3. Scroll down to the **Common Options** section and enable the **Comfort Clients** feature. This will set the option on all of the applicable protocols. The ability to set this feature on a protocol by protocol basis exists in the CLI.
4. Select **OK** or **Apply** to save the changes.
5. Apply this Proxy Options profile in any security policy for it to take effect on all traffic handled by the policy.

The default values for Interval and Amount are 10 and 1, respectively. This means that when client comforting takes effect, 1 byte of the file is sent to the client every 10 seconds. You can change these values to vary the amount and frequency of the data transferred by client comforting.

Oversized files and emails

Downloaded files can range from a few Kilobytes to multiple Gigabytes. A FortiGate doesn't have the memory to allow for a large number of people downloading large files. Imagine the memory required for a team of developers to all download the latest Linux OS distribution at once, in addition to the normal requirements of the firewall. Everything would come to a grinding halt if the FortiGate tried to store each of those Gigabyte+ files in memory. To give you some piece of mind, the chances of malware being in a large file like those is much smaller than in a smaller single Megabyte file, so the threat is somewhat limited, but you will probably want to use your computers antivirus software to scan those large files after they have been downloaded.

A threshold must be set to prevent the resources of the system from becoming overloaded. By default the threshold is 10 MB. Any files larger than the threshold will not be scanned for malware. With a maximum file size threshold in place, it must now be determined what is to be done with the files that are larger than threshold. There are only 2 choices; either the file is passed through without being scanned for malware or the file is blocked. The default action for oversized files is to pass them through.

If you wish to block the downloading of files over the threshold, this can be set within the Proxy Option profile found at **Security Profiles > Proxy Options**, under **Common Options**.

Enable **Block Oversized File/Email**.

This will reveal an additional option, **Threshold (MB)**. The threshold of the files is set based upon the protocol being used to transfer the file. In the CLI and configuration file, the threshold variable is found in each of the protocol sections within the profile. Changing the value in this field will change the `oversize-limit` value for all of the protocols.

If you wish to change the `oversize-limit` value on the protocols covered in a Proxy Option profile you have two options.

1. You can go into the CLI and change the value manually within each of the protocol sections.
2. You can use the GUI to temporarily block oversized files, and when configuring it change the threshold to the new value that you want. Apply this setting. Then go back to the profile and turn off the block setting. If you now go into the CLI you will find that the configuration file has retained the new oversize-limit value.

The settings can be found in the CLI by going to:

```
config firewall profile-protocol-options
  edit <profile_name>
    config <protocol>
      set oversize-limit <size_int>
    end
  end
end
```

Archive scan depth

The antivirus scanner will open archives and scan the files inside. Archives within other archives, or nested archives, are also scanned to a default depth of twelve nestings. You can adjust the number of nested archives the FortiGate unit will scan with the `uncompressed-nest-limit` CLI command. Further, the limit is configured separately for each traffic type.

Configuring archive scan depth

For example, this CLI command sets the archive scan depth for SMTP traffic to 5. That is, archives within archives will be scanned five levels deep.

```
config firewall profile-protocol-options
  edit "default"
    config http
      set uncompressed-nest-limit 5
    end
  end
end
```

You can set the nesting limit from 2 to 100.

Scan buffer size

When checking files for viruses, there is a maximum file size that can be buffered. Files larger than this size are passed without scanning. The default size for all FortiGate models is 10 megabytes.

Archived files are extracted and email attachments are decoded before the FortiGate unit determines if they can fit in the scan buffer. For example, a 7 megabyte ZIP file containing a 12 megabyte EXE file will be passed without scanning with the default buffer size. Although the archive would fit within the buffer, the uncompressed file size will not.

Configuring the uncompression buffer

In this example, the `uncompressed-oversize-limit` CLI command is used to change the scan buffer size to 20 megabytes for files found in HTTP traffic:

```
config firewall profile-protocol-options
  edit <profile_name>
    config http
      set uncompressed-oversize-limit 20
    end
  end
```

```
end
end
```

The maximum buffer size varies by model. Enter `set uncompressed-oversize-limit ?` to display the buffer size range for your FortiGate unit.

Windows file sharing (CIFS)

FortiOS supports virus scanning of Windows file sharing traffic. This includes CIFS, SMB, and SAMBA traffic. This feature is applied by enabling SMB scanning in an antivirus profile and then adding this profile to a security policy that accepts CIFS traffic. CIFS virus scanning is available only through flow-based antivirus scanning.

FortiOS flow-based virus scanning can detect the same number of viruses in CIFS/SMB/SAMBA traffic as it can for all supported content protocols.

Note the following about CIFS/SMB/SAMBA virus scanning:

- Some newer version of SAMBA clients and SMB2 can spread one file across multiple sessions, preventing some viruses from being detected if this occurs.
- Enabling CIFS/SMB/SAMBA virus scanning can affect FortiGate performance.
- SMB2 is a new version of SMB that was first partially implemented in Windows Vista.
- Currently SMB2 is supported by Windows Vista or later, and partly supported by Samba 3.5 and fully support by Samba 3.6.
- The latest version of SMB2.2 will be introduced with Windows 8.
- Most clients still use SMB as default setting.

Configuring CIFS/SMB/SAMBA virus scanning

Use the following command to enable CIFS/SMB/SAMBA virus scanning in an antivirus profile:

```
config antivirus profile
edit <smb-profile>
config smb
set options scan
end
```

Then add this antivirus profile to a security policy that accepts the traffic to be virus scanned. In the security policy the service can be set to ANY, SAMBA, or SMB.

```
config firewall policy
edit <policy-id-integer>
set service ANY
...
set utm-status enable
set av-profile <smb-profile>
end
```

Enabling AntiVirus scanning

Antivirus scanning is configured in an AntiVirus profile, but it is enabled in a firewall policy. Once the use of an AntiVirus profile is enabled and selected in one or more firewall policies, all the traffic controlled by those firewall policies will be scanned according to the settings in that profile.

By going to **System > Feature Visibility**, you can enable or disable two aspects of the AntiVirus Profile.

1. **AntiVirus** will determine if the option to use AntiVirus profiles is available.
2. **Multiple Security Profiles** will determine if you can configure any AntiVirus profiles beyond the default profile.

The use of antivirus protection is a minimum standard for security protection. The question left to decide is whether or not you wish to use multiple profiles in your configuration.

From **Security Profiles > AntiVirus** you can edit existing profiles or create and configure new antivirus profiles that can then be applied to firewall policies. A profile is specific configuration information that defines how the traffic within a firewall policy is examined and what action may be taken based on the examination.

The configuration of the antivirus profile depends on whether the inspection mode is proxy-based or flow-based. You select the inspection mode by going to the **System > Settings** page. The FortiGate's inspection mode is also displayed on the unit's **Dashboard** in the **System Information** widget.

The discussion of the [differences in antivirus scanning modes](#) helps to understand how this scanning works in proxy- and flow-based inspection, as well as in different versions of FortiOS 5.x.

Enabling AntiVirus in Proxy-mode - GUI

1. Go to **Security Profiles > AntiVirus**.
2. Choose whether you want to edit an existing profile or create a new one.
 - The default profile will be the one displayed by default.
 - If you are going to edit an existing profile, selecting it can be done by either using the drop down menu in the upper right hand corner of the window or by selecting the List icon (the furthest right of the 3 icons in the upper right of the window, if resembles a page with some lines on it), and then selecting the profile you want to edit from the list.
 - If you need to create a new profile you can either select the **Create New** icon (a plus sign within a circle) or select the **List icon** and then select the **Create New** link in the upper left of the window that appears.
3. If you are creating a new profile, write a name for it in the **Name** field.
4. For the **Detect Viruses** field, select either **Block** to prevent infected files from passing throughout the FortiGate or **Monitor** to allow infected files to pass through the FortiGate but to record instances of infection.
5. Under **Inspected Protocols**, enable the protocols you wish to be blocked or monitored.
6. Under **APT Protection Options**, you may enable the following: **Content Disarm and Reconstruction**, **Treat Windows Executables in Email Attachments as Viruses** and **Send Files to FortiSandbox Cloud for Inspection**, and **Use Virus Outbreak Prevention Database**.

FortiSandbox options are only available if you have a FortiCloud account active on your FortiGate.

7. Select **Apply**.
8. Add the AntiVirus profile to a firewall security policy.

To view Mobile Malware license and version information, go to **System > FortiGuard** and locate the Mobile Malware section in the **License Information** table.

Content Disarm and Reconstruction (CDR)

Content Disarm and Reconstruction (CDR) is used to remove exploitable content and replace it with content that is known to be safe. As the files are processed through an enabled Proxy-based AntiVirus profile, content that is deemed malicious or unsafe is replaced with content that will allow the traffic to continue, but not put the recipient at risk.

Content that can be scanned includes PDF and Microsoft Office files leaving the network on CDR-supported protocols (HTTP web download, SMTP email send, IMAP/POP3 email retrieval—MAPI is not supported).

This feature will work without FortiSandbox configured, but only if you wish to discard the original file. If FortiSandbox is configured and it responds that the file is clean, it will pass the content unmodified.



This feature will not work if `splice` or `client-comfort` are enabled under `profile-protocol-options` for SMTP.

CDR does not alter documents in an HTTP POST, and is not designed to strip content leaving the network for HTTP. It only works on HTTP GET.

Syntax

The use of CDR is enabled or disabled separately for each protocol in the profile. Note that all CDR commands are only available when you set the profile's `inspection-mode` to `proxy`.

```
config antivirus profile
  edit <name>
    set inspection-mode proxy
    config <protocol>
      set options scan
      set content-disarm {enable | disable}
    next
  end
end
```



You must ensure that `set options scan` is configured.

If `set options av-monitor` is configured for a protocol, it will enable the `detect-only` option (see below) and CDR will not occur for that protocol.

The enabling and disabling of the CDR is specific to the protocol, but the granular configuration of which types of content will be rewritten by the CDR engine are configured based on the AntiVirus profile. The settings within the `config content-disarm` context are applicable to all of the CDR enabled protocols.

```
config antivirus profile
  edit <name>
    config content-disarm
      set original-file-destination {fortisandbox | quarantine | discard}
      set office-macro {enable | disable}
      set office-hylink {enable | disable}
      set office-linked {enable | disable}
      set office-embed {enable | disable}
      set pdf-javacode {enable | disable}
      set pdf-embedfile {enable | disable}
      set pdf-act-gotor {enable | disable}
      set pdf-act-launch {enable | disable}
      set pdf-act-uri {enable | disable}
      set pdf-act-sound {enable | disable}
      set pdf-act-movie {enable | disable}
      set pdf-act-java {enable | disable}
      set pdf-act-form {enable | disable}
      set cover-page {enable | disable}
      set detect-only {enable | disable}
    next
  end
end
```

Where:

Option	Description
<code>original-file-destination</code>	Select the destination to which files will be sent for inspection. Note that, once you enable <code>content-disarm</code> under a protocol, you will be warned that all original files will be discarded. To be able to retrieve the original files, you must set an <code>original-file-destination</code> for this profile.
<code>office-macro</code>	Enables/disables stripping of macros in Microsoft Office documents.
<code>office-hylink</code>	Enables/disables stripping of hyperlinks in Microsoft Office documents.
<code>office-linked</code>	Enables/disables stripping of linked objects in Microsoft Office documents.
<code>office-embed</code>	Enables/disables stripping of embedded objects in Microsoft Office documents.
<code>pdf-javacode</code>	Enables/disables stripping of JavaScript code in PDF documents.
<code>pdf-embedfile</code>	Enables/disables stripping of embedded files in PDF documents.
<code>pdf-act-gotor</code>	Enables/disables stripping of links to other PDFs in PDF documents.
<code>pdf-act-launch</code>	Enables/disables stripping of links to external applications in PDF documents.
<code>pdf-act-uri</code>	Enables/disables stripping of links to URI resources in PDF documents.
<code>pdf-act-sound</code>	Enables/disables stripping of embedded sound files in PDF documents.
<code>pdf-act-movie</code>	Enables/disables stripping of embedded movies in PDF documents.
<code>pdf-act-java</code>	Enables/disables stripping of actions that execute JavaScript code in PDF documents.
<code>pdf-act-form</code>	Enables/disables stripping of actions that submit data to other targets in PDF documents.
<code>cover-page</code>	Enables/disables inserting a cover page into the disarmed document.
<code>detect-only</code>	Enables/disables only detect disarmable files, do not alter content.

When the antivirus profile successfully detects suspicious content and strips the data, a new page is appended to the start of the document with a message that reads *"This file has been cleaned of potential threats"*.

You can set `cover-page disable` (see above) if you do not want a cover page appended to any disarmed content.

FortiGuard virus outbreak prevention

FortiGuard virus outbreak prevention uses checksums to filter files in order to detect and prevent quick virus outbreaks, because it usually takes at least a few hours for FortiGuard to develop and push signatures and a virus outbreak can do a lot of damage within that time period. This method proves to be quite effective using hash values of probable virus files.

Enable this feature under **Security Profiles > AntiVirus > Use Virus Outbreak Prevention Database**. Note that this feature requires a license, which you can obtain through **System > FortiGuard > Outbreak Prevention**.

Syntax

Note that `outbreak-prevention` is only available when `options` is set to `scan`:

```
config antivirus profile
  edit <name>
    config <protocol>
      set options scan
      set outbreak-prevention {disabled | files | full-archive}
    next
  ...
```

where `full-archive` analyzes files including the contents of archives, as opposed to `files` which does not include the contents of archives.

Enabling AntiVirus in Flow-mode - GUI

1. Go to **Security Profiles > AntiVirus**.
2. Choose whether you want to edit an existing profile or create a new one.
 - The default profile will be the one displayed by default.
 - If you are going to edit an existing profile, selecting it can be done by either using the drop down menu in the upper right hand corner of the window or by selecting the List icon (the furthest right of the 3 icons in the upper right of the window, if resembles a page with some lines on it), and then selecting the profile you want to edit from the list.
 - If you need to create a new profile you can either select the **Create New** icon (a plus sign within a circle) or select the **List icon** and then select the **Create New** link in the upper left of the window that appears.
3. If you are creating a new profile, write a name for it in the **Name** field.
4. Select **Quick** or **Full Scan Mode**(see the discussion of the [differences in antivirus scanning modes](#) for more information).
5. For the **Detect Viruses** field, select either **Block** to prevent infected files from passing throughout the FortiGate or **Monitor** to allow infected files to pass through the FortiGate but to record instances of infection.
6. Under **Inspected Protocols**, enable the protocols you wish to be blocked or monitored.
7. Under **Inspection Options**, you may enable the following: **Treat Windows Executables in Email Attachments as Viruses** and **Include Mobile Malware Protection**.



You may also enable the following options if you have a FortiCloud account active on your FortiGate: **Send Files to FortiSandbox Cloud for Inspection** and **Use FortiSandbox Database**.

8. Select **OK** or **Apply**.
9. Add the AntiVirus profile to a firewall security policy.

Enabling AntiVirus - CLI

Configure the scan option for each type of traffic you want scanned.

1. Configure the AntiVirus profile

```
config antivirus profile
  edit <profile_name>
    set comment "scan and delete virus"
    set replacemsg-group ''
    set scan-botnet-connections block
    set ftgd-analytics suspicious
    config http
      set options scan
    end
    config ftp
      set options scan
    end
    config imap
      set options scan
    end
    config pop3
      set options scan
    end
    config smtp
      set options scan
    end
    config nntp
      set options scan
    end
    config smb
      set options scan
    end
  end
end
```

2. Add the AntiVirus profile to the Fortigate firewall security policy. When using the CLI, you will need to know the policy ID number.

```
config firewall policy
  edit <policy ID number>
    set av-profile <profile_name>
    set profile-protocol-options default
  end
end
```

Overriding the AV engine file scan timeout

Overriding the AV engine file scan timeout allows the FortiGate to scan files as large as 4GB without breaking the scan.

Override the large file scan timeout value in seconds (30 - 3600). Zero is the default value and is used to disable this command. When disabled, the daemon adjusts the large file scan timeout based on the file size.

Syntax

```
config antivirus settings
```



```
set override-timeout 0
end
```

Testing your antivirus configuration

You have configured your FortiGate unit to stop viruses, but you'd like to confirm your settings are correct. Even if you have a real virus, it would be dangerous to use for this purpose. An incorrect configuration will allow the virus to infect your network.

To solve this problem, the European Institute of Computer Anti-virus Research has developed a test file that allows you to test your antivirus configuration. The EICAR test file is not a virus. It cannot infect computers, nor can it spread or cause any damage. It's a very small file that contains a sequence of characters. Your FortiGate unit recognizes the EICAR test file as a virus so you can safely test your FortiGate unit antivirus configuration.

Go to <http://www.fortiguard.com/antivirus/eicartest.html> to download the test file (eicar.com) or the test file in a ZIP archive (eicar.zip).

If the antivirus profile applied to the security policy that allows you access to the Web is configured to scan HTTP traffic for viruses, any attempt to download the test file will be blocked. This indicates that you are protected.

Example scenarios

The following examples provide sample antivirus configuration scenarios.

Configuring simple default AntiVirus profile

If performance is not a real concern and the FortiGate's resources are not being stretched, it is perfectly reasonable to create one AntiVirus profile that covers the range of uses found in your environment. This example is one possible default configuration.

Context:

- This is an edited default profile and will be used on all security policies
- It will need to scan for malware on all available protocols.
- Malware, botnets, and grayware should be blocked
- The inspection method should be flow-based
- A current FortiCloud account is available

Creating the profile - GUI

Edit AntiVirus Profile default ▼

Name

Comments 29/255

Scan Mode Quick Full

Detect Viruses Block Monitor

APT Protection Options

Treat Windows Executables in Email Attachments as Viruses ☒

Send Files to FortiSandbox Appliance for Inspection None All Supported Files

Do not submit files matching types

Do not submit files matching file name patterns

Use Virus Outbreak Prevention Database i ☒

Use FortiSandbox Database i ☐

Apply

1. In the following fields, enter the settings shown in the screenshot.
2. Select **Apply**.
3. Enable grayware scanning through the CLI.


```
config antivirus settings
  set grayware enable
end
```

Creating the profile - CLI

1. Enter the CLI by one of the following methods:
 - SSH through a terminal emulator
 - CLI Console access
 - FortiExplorer's CLI mode
2. Enter the following commands:


```
config antivirus profile
  edit default
    set comment "scan and delete virus"
    set inspection-mode flow-based
    set scan-botnet-connections block
    set ftgd-analytics suspicious
  config http
    set options scan
  end
  config ftp
    set options scan
  end
```

```
config imap
    set options scan
end
config pop3
    set options scan
end
config smtp
    set options scan
end
config nntp
    set options scan
end
config smb
    set options scan
end
end
```

3. Enable grayware scanning

```
config antivirus settings
    set grayware enable
end
```

Setting up a basic proxy-based AntiVirus profile for email traffic

Small offices, whether they are small companies, home offices, or satellite offices, often have very simple needs. This example details how to enable antivirus protection on a FortiGate unit located in a satellite office.

Context:

- The satellite office does not have an internal email server. To send and retrieve email, the employees connect to an external mail server.
- There is a specific firewall security profile that handles the email traffic from the Internet to the mail server. The only traffic on this policy will be POP3 and IMAP and SMTP
- The company policy is to block viruses and connections to botnets.
- The FortiGate unit is a small model and the Internet bandwidth is limited so the policy is to not submit files to the FortiSandbox.

Creating the profile - GUI

Edit AntiVirus Profile
default

Name
default

Comments
Scan files and block viruses.
29/255

Detect Viruses
Block Monitor

Inspected Protocols

HTTP
SMTP
POP3
IMAP
MAPI
FTP

APT Protection Options

Content Disarm and Reconstruction
Treat Windows Executables in Email Attachments as Viruses
Send Files to FortiSandbox Appliance for Inspection
None All Supported Files
Do not submit files matching types
Do not submit files matching file name patterns
Use Virus Outbreak Prevention Database
Use FortiSandbox Database

Apply

1. In the following fields, enter the settings shown below:

Name	email-av
Comments	Scans email traffic from Internet for malware
Detect Viruses	Block
Inspected Protocols	all checked (HTTP, SMTP, POP3, IMAP, MAPI, and FTP).
Content Disarm and Reconstruction	checked (optional) - used to remove exploitable content and replace it with content that is known to be safe. For more information, see Content Disarm and Reconstruction (CDR)

Original File Destination	Destination to which files will be sent for inspection: FortiSandbox, File Quarantine, or Discard.
Treat Windows Executables in Email Attachments as Viruses	checked - also optionally decide whether or not to submit files matching particular types and/or file name patterns.
Send Files to FortiSandbox Appliance for Inspection	checked (All Supported Files).
Use Virus Outbreak Prevention Database	checked - used to preempt outbreaks before AV Signatures are created.
Use FortiSandbox Database	checked - supplements the AV Signature database.

2. Select **Apply**.

Creating the profile - CLI

1. Enter the CLI by one of the following methods:
 - SSH through a terminal emulator
 - CLI Console widget
 - FortiExplorer's CLI mode

2. Enter the following commands:

```
config antivirus profile
  edit "email-av"
    set comment "Scans email traffic from Internet for malware"
    set inspection-mode proxy
    config content-disarm
      set original-file-destination {fortisandbox | quarantine | discard}
      set ...
    config <protocol>
      set options scan
    end
  end
end
```

3. Additionally, if you wish to only send those files to FortiSandbox that heuristics determines as suspicious, enter the following (only available via the CLI):

```
config antivirus profile
  edit "email-av"
    set ftgd-analytics suspicious
  end
```

For more information on how to strip content from various content types from documents (hyperlinks, linked objects, embedded objects, JavaScript code), see [Content Disarm and Reconstruction \(CDR\)](#) and the [FortiOS 6.0 CLI Reference](#).

Adding the profile to a policy

In this scenario the following assumptions will be made:

- The policy that the profile is going to be added to is an IPv4 policy.
- The ID number of the policy is 11.

- The AntiVirus profile being added will be the "default" profile
- The SSL/SSH Inspection profile used will be the "default" profile



FortiClient enforcement has been moved from the Policy page to **Network > Interfaces** to enforce FortiClient registration on a desired LAN interface rather than a policy.

Adding the profile - GUI

1. Go to **Policy & Objects > IPv4 Policy**.
2. Use your preferred method of finding a policy.
 - If the ID column is available you can use that.
 - You can also choose based on your knowledge of the parameters of the policy
 - Select the policy with ID value of 11
3. In the Edit Policy window, go to the Security Profiles section
4. Turn ON AntiVirus, and in the drop down menu for the field, select default
5. If the AntiVirus profile is proxy-based the Proxy Options field and drop down menu will be revealed.
6. The SSL/SSH Inspection field will automatically be set to ON and one of the profiles will need to be selected from the drop down menu. In this case default is selected.
7. The log options will depend on your requirements and resources but to verify that everything is working properly, it is a good idea to turn ON logging of All Sessions after setting up a new profile and after giving some time for logs to accumulate
8. Turn on Antivirus.
9. Select an antivirus profile.
10. Select **OK** to save the security policy.

Adding the profile - CLI

To select the antivirus profile in a security policy — CLI

```
config firewall policy
  edit 11
    set utm-status enable
    set profile-protocol-options default
    set av-profile basic_antivirus
  end
```

Block files larger than 8 MB

Set proxy options profile to block files larger than 8 MB

1. Go to **Security Profiles > Proxy Options**.
2. Edit the default or select Create New to add a new one.
3. Scroll down to the common Options Section and place a check in the box next to BlockOversized File/Email
4. The sub line Threshold (MB) will appear with a value field. Enter 8.
5. Select **OK** or **Apply**.

The proxy options profile is configured, but to block files, you must select it in the firewall policies handling the traffic that contains the files you want blocked.

To select the Proxy Options profile in a security policy

1. Go to **Policy & Objects > IPv4 Policy** (or **IPv6 Policy**, depending).
2. Edit or create a security policy.
3. Select a proxy-based security profile. You will know that there is a proxy component to the Security Profile because when a Security Profile is Proxy based the Proxy Options field will be visible (for example, select an Antivirus profile that includes proxy scanning).
4. Beside Proxy Options select the name of the MTU proxy options protocol.
5. Select **OK** to save the security policy.
6. Once you complete these steps, any files in the traffic subject to Security Profile scanning handled by this policy that are larger than 8MB will be blocked. If you have multiple firewall policies, examine each to determine if you want to apply similar file blocking the them as well.

Web filter

This section describes FortiGate web filtering for HTTP traffic. The three main parts of the web filtering function, the Web Content Filter, the URL Filter, and the FortiGuard Web Filtering Service interact with each other to provide maximum control over what users on your network can view as well as protection to your network from many Internet content threats. Web Content Filter blocks web pages containing words or patterns that you specify. URL filtering uses URLs and URL patterns to block or exempt web pages from specific sources. FortiGuard Web Filtering provides many additional categories you can use to filter web traffic.

This Handbook chapter includes [Inside FortiOS: Web Filtering](#) and provides readers an overview of the features and benefits of key FortiOS components.

For further detail than the Inside FortiOS document, we provide the following topics:

Web filter concepts

Web filtering is a means of controlling the content that an Internet user is able to view. With the popularity of web applications, the need to monitor and control web access is becoming a key component of secure content management systems that employ antivirus, web filtering, and messaging security. Important reasons for controlling web content include:

- lost productivity because employees are accessing the web for non-business reasons
- network congestion — when valuable bandwidth is used for non-business purposes, legitimate business applications suffer
- loss or exposure of confidential information through chat sites, non-approved email systems, instant messaging, and peer-to-peer file sharing
- increased exposure to web-based threats as employees surf non-business-related web sites
- legal liability when employees access/download inappropriate and offensive material
- copyright infringement caused by employees downloading and/or distributing copyrighted material.

As the number and severity of threats increase on the World Wide Web, the risk potential increases within a company's network as well. Casual non-business related web surfing has caused many businesses countless hours of legal litigation as hostile environments have been created by employees who download and view offensive content. Web-based attacks and threats are also becoming increasingly sophisticated. Threats and web-based applications that cause additional problems for corporations include:

- spyware/grayware
- phishing
- pharming
- instant messaging
- peer-to-peer file sharing
- streaming media
- blended network attacks.

Spyware, also known as grayware, is a type of computer program that attaches itself to a user's operating system. It does this without the user's consent or knowledge. It usually ends up on a computer because of

something the user does such as clicking on a button in a pop-up window. Spyware can track the user's Internet usage, cause unwanted pop-up windows, and even direct the user to a host web site. For further information, visit the FortiGuard Center.

Some of the most common types of grayware infection occur when:

- downloading shareware, freeware, or other forms of file-sharing services
- clicking on pop-up advertising
- visiting legitimate web sites infected with grayware.

Phishing is the term used to describe attacks that use web technology to trick users into revealing personal or financial information. Phishing attacks use web sites and email that claim to be from legitimate financial institutions to trick the viewer into believing that they are legitimate. Although phishing is initiated by spam email, getting the user to access the attacker's web site is always the next step.

Pharming is a next generation threat that is designed to identify and extract financial, and other key pieces of information for identity theft. Pharming is much more dangerous than phishing because it is designed to be completely hidden from the end user. Unlike phishing attacks that send out spam email requiring the user to click to a fraudulent URL, pharming attacks require no action from the user outside of their regular web surfing activities. Pharming attacks succeed by redirecting users from legitimate web sites to similar fraudulent web sites that have been created to look and feel like the authentic web site.

Instant messaging presents a number of problems. Instant messaging can be used to infect computers with spyware and viruses. Phishing attacks can be made using instant messaging. There is also a danger that employees may use instant messaging to release sensitive information to an outsider.

Peer-to-peer (P2P) networks are used for file sharing. Such files may contain viruses. Peer-to-peer applications take up valuable network resources and may lower employee productivity but also have legal implications with the downloading of copyrighted or sensitive company material.

Streaming media is a method of delivering multimedia, usually in the form of audio or video to Internet users. Viewing streaming media impacts legitimate business by using valuable bandwidth.

Blended network threats are rising and the sophistication of network threats is increasing with each new attack. Attackers learn from each successful attack and enhance and update their attack code to become more dangerous and to spread faster. Blended attacks use a combination of methods to spread and cause damage. Using virus or network worm techniques combined with known system vulnerabilities, blended threats can quickly spread through email, web sites, and Trojan applications. Examples of blended threats include Nimda, Code Red, Slammer, and Blaster. Blended attacks can be designed to perform different types of attacks, which include disrupting network services, destroying or stealing information, and installing stealthy backdoor applications to grant remote access.

Different ways of controlling access

The methods available for monitoring and controlling Internet access range from manual and educational methods to fully automated systems designed to scan, inspect, rate and control web activity.

Common web access control mechanisms include:

- establishing and implementing a well-written usage policy in the organization on proper Internet, email, and computer conduct
- installing monitoring tools that record and report on Internet usage
- implementing policy-based tools that capture, rate, and block URLs.

The following information shows how the filters interact and how to use them to your advantage.

Order of web filtering

The FortiGate unit applies web filters in a specific order:

1. URL filter
2. FortiGuard Web Filter
3. web content filter
4. web script filter
5. antivirus scanning.

If you have blocked a FortiGuard Web Filter category but want certain users to have access to URLs within that pattern, you can use the **Override** within the FortiGuard Web Filter. This will allow you to specify which users have access to which blocked URLs and how long they have that access. For example, if you want a user to be able to access www.example.com for one hour, you can use the override to set up the exemption. Any user listed in an override must fill out an online authentication form that is presented when they try to access a blocked URL before the FortiGate unit will grant access to it.

If you have blocked a FortiGuard Web Filter category but want users within a specific Web Filter profile to have access to URLs within that pattern, you can use the following CLI command below to override (this will have no timeout affiliated to it):

CLI syntax:

```
config webfilter profile
  edit <profile>
    config web
      set whitelist exempt-av exempt-dlp exempt-rangeblock extended-log-others
    end
  end
```

This command will set a Web Filter profile that exempts AV, DLP, RangeBlock, and supports extended log by FortiGuard whitelist.

Inspection modes

This topic briefly discusses proxy and flow-based inspection modes. For more information on flow vs. proxy inspection modes on your FortiGate and how they impact web filtering, see [Individual Security Profile considerations](#) in the [Inspection Modes](#) section.

Proxy

Proxy-based inspection involves buffering traffic and examining it as a whole before determining an action. The process of having the whole of the data to analyze allows for the examination of more points of data than the flow-based or [DNS methods](#).

The advantage of a proxy-based method is that the inspection can be more thorough than the other methods, yielding fewer false positive or negative results in the data analysis.

Flow-based

The flow-based inspection method examines the file as it passes through the FortiGate unit without any buffering. As each packet of the traffic arrives it is processed and forwarded without waiting for the complete file or web page.

The advantage of the flow-based method is that the user sees a faster response time for HTTP requests and there is less chance of a time-out error due to the server at the other end responding slowly.

The disadvantages of this method are: (1) there is a higher probability of a false positive or negative in the analysis of the data; and, (2) a number of security features that can be used in the proxy-based method are not available in the flow-based inspection method. There are also fewer actions available based on the categorization of the website by FortiGuard services.

In flow mode, Web Filter profiles only include flow-mode features. Web filtering is still done with the same engines and to the same accuracy, but some inspection options are limited or not available in flow mode.

Configuring Web Filter profiles in flow-mode is different depending on the **NGFW mode** selected.



See ["What's new in FortiOS 6.0.1" on page 2381](#) and [Individual Security Profile Considerations](#) in the [Inspection Modes](#) section for more details on flow vs. proxy inspection modes on your FortiGate.

FortiGuard Web Filtering Service

FortiGuard Web Filtering is a managed web filtering solution available by subscription from Fortinet. Before you begin to use the FortiGuard Web Filtering options, verify that you have a valid subscription to the service for your FortiGate firewall.

FortiGuard Web Filtering enhances the web filtering features supplied with your FortiGate unit by sorting billions of web pages into a wide range of categories users can allow or block. The FortiGate unit accesses the nearest FortiGuard Web Filtering Service Point to determine the category of a requested web page, and then applies the security policy configured for that user or interface. FortiGuard Web Filtering supports detection for traffic using HTTP protocol (versions 1.0, 1.1, and 2.0).

FortiGuard Web Filtering includes over 45 million individual ratings of web sites that apply to more than two billion pages. Pages are sorted and rated into several dozen categories administrators can allow or block. Categories may be added or updated as the Internet evolves. To make configuration simpler, you can also choose to allow or block entire groups of categories. Blocked pages are replaced with a message indicating that the page is not accessible according to the Internet usage policy.

FortiGuard Web Filtering ratings are performed by a combination of proprietary methods including text analysis, exploitation of the web structure, and human raters. Users can notify the FortiGuard Web Filtering Service Points if they feel a web page is not categorized correctly, so that the service can update the categories in a timely fashion.

FortiGuard web filtering and your FortiGate unit

When FortiGuard Web Filtering is enabled in a web filter or a DNS filter profile, the setting is applied to all firewall policies that use this profile. When a request for a web page appears in traffic controlled by one of these firewall policies, the URL is sent to the nearest FortiGuard server. The URL category is returned. If the category is blocked, the FortiGate unit provides a replacement message in place of the requested page. If the category is not blocked, the page request is sent to the requested URL as normal.

FortiGuard web filtering actions

The possible actions are:

- **Allow** permits access to the sites within the category.
- **Block** prevents access to sites within the category. Users attempting to access a blocked site will receive a replacement message explaining that access to the site is blocked.
- **Monitor** permits and logs access to sites in the category. You may also enable user quotas when enabling the monitor action.
- **Warning** presents the user with a message, allowing them to continue if they choose.
- **Authenticate** requires a user to authenticate with the FortiGate unit before being allowed access to the category or category group.


The options of actions available will depend on the mode of inspection.

- Proxy - Allow, Block, Monitor, Warning, Authenticate and Disable.
- Flow-based, policy-based - Allow, Block & Monitor.
- Flow-based, profile-based - Allow, Deny



Configuring Web Filter profiles in flow-mode is different depending on the [NGFW mode](#) selected.

Web filtering flowchart



FortiGuard web filtering categories

The following tables identify each FortiGuard web filtering category (organized by group) along with associated category IDs. You can access the current list of category IDs through the CLI.

```
config webfilter profile
edit default
config ftgd-wf
config filters
edit 1
set category ?
```

For a complete description of each web filtering category, visit <http://www.fortiguards.com/webfilter/categories>.

Potentially Liable

ID	Category
1	Drug Abuse
3	Hacking
4	Illegal or Unethical
5	Discrimination
6	Explicit Violence

ID	Category
12	Extremist Groups
59	Proxy Avoidance
62	Plagiarism
83	Child Abuse

Adult/Mature Content

ID	Category
2	Alternative Beliefs
7	Abortion
8	Other Adult Materials
9	Advocacy Organizations
11	Gambling
13	Nudity and Risque
14	Pornography
15	Dating

ID	Category
16	Weapons (Sales)
57	Marijuana
63	Sex Education
64	Alcohol
65	Tobacco
66	Lingerie and Swimsuit
67	Sports Hunting and War Games

Bandwidth Consuming

ID	Category
19	Freeware and Software Downloads
24	File Sharing and Storage
25	Streaming Media and Download

ID	Category
72	Peer-to-peer File Sharing
75	Internet Radio and TV
76	Internet Telephony

Security Risk

ID	Category
26	Malicious Websites
61	Phishing
	Newly Observed Domain

ID	Category
86	Spam URLs
88	Dynamic DNS
	Newly Registered Domain

Newly observed domain (NOD) applies to URLs whose domain name is not rated and were observed for the first time in the past 30 minutes.

Newly registered domain (NRD) applies to URLs whose domain name was registered in the previous 10 days.

General Interest - Personal

ID	Category
17	Advertising
18	Brokerage and Trading
20	Games
23	Web-based Email
28	Entertainment
29	Arts and Culture
30	Education
33	Health and Wellness
34	Job Search
35	Medicine

ID	Category
47	Travel
48	Personal Vehicles
54	Dynamic Content
55	Meaningless Content
58	Folklore
68	Web Chat
69	Instant Messaging
70	Newsgroups and Message Boards
71	Digital Postcards
77	Child Education

ID	Category
36	News and Media
37	Social Networking
38	Political Organizations
39	Reference
40	Global Religion
42	Shopping
44	Society and Lifestyles
46	Sports

ID	Category
78	Real Estate
79	Restaurant and Dining
80	Personal Websites and Blogs
82	Content Servers
85	Domain Parking
87	Personal Privacy
89	Auction

General Interest - Business

ID	Category
31	Finance and Banking
41	Search Engines and Portals
43	General Organizations
49	Business
50	Information and Computer Security
51	Government and Legal Organizations

ID	Category
52	Information Technology
53	Armed Forces
56	Web Hosting
81	Secure Websites
84	Web-based Applications

Local categories

Users can define custom or local categories. See [Overriding FortiGuard Website Categorization](#) for details.

FortiGuard web filtering usage quotas

In addition to using category and classification blocks and overrides to limit user access to URLs, you can set a daily quota by category, category group, or classification. Quotas allow access for a specified length of time or a specific bandwidth, calculated separately for each user. Quotas are reset every day at midnight.

Users must authenticate with the FortiGate unit. The quota is applied to each user individually so the FortiGate must be able to identify each user. One way to do this is to configure a security policy using the identity-based policy feature. Apply the web filter profile in which you have configured FortiGuard Web Filter and FortiGuard Web Filter quotas to such a security policy.



The use of FortiGuard Web Filtering quotas requires that users authenticate to gain web access. The quotas are ignored if applied to a security policy in which user authentication is not required.

Editing the web filter profile resets the quota timers for all users.

When a user first attempts to access a URL, they're prompted to authenticate with the FortiGate unit. When they provide their user name and password, the FortiGate unit recognizes them, determines their quota allowances, and monitors their web use. The category and classification of each page they visit is checked and FortiGate unit adjusts the user's remaining available quota for the category or classification.

Quota hierarchy

You can apply quotas to categories and category groups. Only one quota per user can be active at any one time. The one used depends on how you configure the FortiGuard Web Filter.

When a user visits a URL, the FortiGate unit queries the FortiGuard servers for the category of the URL. From highest to lowest, the relative priority of the quotas are:

1. Category
2. Category group

Configuring web filter profiles

Enabling FortiGuard web filter

FortiGuard Web Filter is enabled and configured within web filter profiles by enabling FortiGuard Categories. The service is engaged by turning on the Web Filter profile and selecting a profile that has FortiGuard Categories enabled on one or more active policies being run by the firewall.

There is also a system wide setting for the enabling or disabling of FortiGuard Web Filter that is only in the CLI.

```
config system fortiguard
    set webfilter-force-off
```

The two options on this setting are enable or disable. The syntax of the settings name is "force-off" so in order to enable FortiGuard Webfilter you have to choose disable for the setting and enable if you want to turn it off.

General configuration steps

1. Go to **Security Profiles > Web Filter**.
2. Determine if you wish to create a new profile, edit an existing one, or clone and edit an existing one.
3. If you are using FortiGuard Categories, enable the FortiGuard Categories, select the categories and select the action to be performed.
4. Configure any **Category Usage Quotas** needed. (Proxy Mode)
5. Allow blocked override if required. (Proxy Mode)
6. Set up **Safe Search** settings and/or YouTube Education settings. (Proxy & Flow-based)
7. Configure **Static URL Settings**. (All Modes)
8. Configure **Rating Options**. (All Modes)
9. Configure **Proxy Options**.
10. Save the filter and web filter profile.

11. To complete the configuration, you need to select the security policy controlling the network traffic you want to restrict. Then, in the security policy, enable Web Filter and select the appropriate web filter profile from the list.

Configuring FortiGuard Web Filter settings

FortiGuard Web Filter includes a number of settings that allow you to determine various aspects of the filtering behavior.

Getting to the Edit Web Filter Profile configuration window

Once you have gotten to the profile configuration window there are a number of settings that can be used, most of which are optional. We will treat each of these options separately, but present the common instructions of how to get to the profile editing page here.

1. Go to **Security Profiles > Web Filter**.
2. Determine if you wish to create a new profile, edit an existing one, or clone and then edit an existing one.
 - a. New profile:
 - i. Select the **Create New** icon, in the upper right of the window (looks like a plus sign in a circle) **OR**
 - ii. Select the **List** icon, in the upper right (looks like a white rectangle with lines like text). Select the **Create New** icon in the upper left.
 - b. Edit existing profile:
 - i. Select the name of the profile that you wish to edit from the drop-down menu **OR**
 - ii. Select the **List** icon, in the upper right (looks like a white rectangle with lines like text). Highlight the name of the profile from the list and select **Edit** from the options above the list.
 - c. Clone a profile:
 - i. Select **Clone** icon in the upper right corner of the window (looks like one square overlapping another) **OR**
 - ii. Select the **List** icon, in the upper right (looks like a white rectangle with lines like text). Highlight the name of the profile from the list and select **Clone** from the options above the list.
3. Make sure there is a valid name, and comment if you want.
4. Configure the settings to best achieve your specific requirements
5. Select **Apply** or **OK**, depending on whether you are editing, creating, or cloning a profile.



In older versions of FortiOS there was a character limitation for the URL of 2048 bytes or approximately 321 characters. If the URL you were trying to reach was longer the URL sent to FortiGuard would be truncated and the service would be unable to categorize the site. Starting in version 5 of the firmware, the parsed URL has been increase to 4 Kilobytes, effectively doubling the length of a URL capable of being categorized.

To configure the FortiGuard Web Filter categories

1. Go to the **Edit Web Filter Profile** window.
2. The category groups are listed in a widget. You can expand each category group to view and configure every sub-category individually within the groups. If you change the setting of a category group, all categories within the group inherit the change.
3. Select the category groups and categories to which you want to apply an action.
To assign an action to a category left click on the category and select from the pop up menu.
4. Select **Apply** or **OK**.

Apply the web filter profile to an identity-based security policy. All the users subject to that policy are restricted by the quotas.



If you look at your logs carefully, you may notice that not every URL connection in the log shows a category. They are left blank. If you take one of those URL and enter it in the FortiGuard website designed to show the category for a URL it will successfully categorize it.

The reason for this is that to optimize speed throughput and reduce the load on the FortiGuard servers the FortiGate does not determine a category rating on scripts and css files.

Configuring FortiGuard Category quotas

1. Go to the **Edit Web Filter Profile** window
2. Verify that the categories that need to have quotas on them are set to one of these actions:
 - **Monitor**
 - **Warning**
 - **Authenticate**
3. Under **Category Usage Quota**, Select **Create New or Edit**
4. In the **New/Edit Quota** window that pops up, enable or disable the specific categories for that quota.
5. At the bottom of the widget, select a quota type and daily allowance for each user:
 - **Time** -- can be entered in **Hours**, **Minutes**, or **Seconds**.
 - **Traffic** -- can be entered in **Bytes**, **KB**, **MB**, or **GB** The value must be greater than 0.
6. Select **Apply or OK**.
7. Continue with any other configuration in the profile
8. Select **Apply or OK**.

Apply the web filter profile to an identity-based security policy. All the users subject to that policy are restricted by the quotas.



The use of FortiGuard Web Filtering quotas requires that users authenticate to gain web access. The quotas are ignored if applied to a security policy in which user authentication is not required.

Editing the web filter profile resets the quota timers for all users.

Configure Allowed Blocked Overrides

1. Go to the **Edit Web Filter Profile** window.
2. Enable **Allow Blocked Override**
3. In the Apply to Group(s) field select the desired **User Group**
4. In the Assign to Profile field, select the desired profile

Configure search engine

There are 2 primary configuration settings in this section.

Enable SafeSearch

To enable the SafeSearch settings

1. Go to the **Edit Web Filter Profile** window.
2. Enable **SafeSearch**
3. Enable Search Engine SafeSearch
4. Enable YouTube Filter
 - a. Enter the YouTube User ID in the Text field



Web Filter in flow mode does not support Safe Search

Log all search keywords

In the GUI, the configuration setting is limited to a checkbox.

Configure static URL filter

Web content filter

To enable the web content filter and set the content block threshold

1. Go to the **Edit Web Filter Profile** window.
2. In the **Static URL Filter** section enable **Web Content Filter**.
3. Select **Create New**.
4. Select the **Pattern Type**.
5. Enter the content **Pattern**.
6. Enter the **Language** from the dropdown menu.
7. Select **Block** or **Exempt**, as required, from the **Action** list.
8. Select **Enable**.
9. Select **OK**.

Configure rating options

Allow Websites When a Rating error Occurs

In the GUI, the configuration setting is limited to a checkbox.

Rate URLs by Domain and IP Address

In the GUI, the configuration setting is limited to a checkbox.

Block HTTP Redirects by Rating

In the GUI, the configuration setting is limited to a checkbox.

Rate Images by URL (Blocked images will be replaced with blanks)

In the GUI, the configuration setting is limited to a checkbox.

Configure Proxy Options

Restrict Google Account Usage to Specific Domains

Configuring the feature in the GUI

Go to **Security Profiles > Web Filter**.

In the **Proxy Options** section, check the box next to **Restrict to Corporate Google Accounts Only**.

Use the **Create New** link within the widget to add the appropriate Google domains that will be allowed.

Configuring the feature in the CLI

To configure this option in the CLI, the URL filter must refer to a web-proxy profile that is using the Modifying HTTP Request Headers feature. The command is only visible when the action for the entry in the URL filter is set to either allow or monitor.

1. Configure the proxy options:

```
config web-proxy profile
  edit "googleproxy"
    config headers
      edit 1
        set name "X-GoogApps-Allowed-Domains"
        set content "fortinet.com, Ladan.ca"
      end
    end
  end
end
```

2. Set a web filter profile to use the proxy options

```
config webfilter urlfilter
  edit 1
    config entries
      edit "*.google.com"
        set type wildcard
        set action {allow | monitor}
        set web-proxy-profile <profile>
      end
    end
  end
end
```

In the CLI, you can also add, modify, and remove header fields in HTTP request when scanning web traffic in proxy-mode. If a header field exists when your FortiGate receives the request, its content will be modified based on the configurations in the URL filter.

Web Resume Download block

In the GUI, the configuration setting is limited to a checkbox.

Provide Details for Blocked HTTP 4xx and 5xx Errors

In the GUI, the configuration setting is limited to a checkbox.

HTTP POST Action

Remove Java Applet Filter

In the GUI, the configuration setting is limited to a checkbox.

Remove ActiveX Filter

In the GUI, the configuration setting is limited to a checkbox.

Remove Cookie Filter

In the GUI, the configuration setting is limited to a checkbox.

Overriding FortiGuard website categorization

In most things there is an exception to the rule. When it comes to the rules about who is allowed to go to which websites in spite of the rules or in this case, policies, it seems that there are more exceptions than to most rules. There are numerous valid reasons and scenarios for exceptions so it follows that there needs to be a way to accommodate this exception.

The different methods of override

There are two different ways to override web filtering behavior based on FortiGuard categorization of a websites if you are operating in proxy-based inspection.

The second method has two variations in implementation and each of the three has a different level of granularity.

1. Using Alternate Categories

Web Rating Overrides

This method manually assigns a specific website to a different Fortinet category or a locally created category.

2. Using Alternate Profiles

Administrative Override or Allow users to override blocked categories

In this method all of the traffic going through the FortiGate unit, using identity based policies and a Web Filtering profile has the option where configured users or IP addresses can use an alternative Web Filter profile when attempting to access blocked websites.

Using Alternate Categories

Web Rating Overrides

There are two approaches to overriding the FortiGuard Web Filtering. The first is an identity-based method that can be configured using a combination of identity-based policies and specifically designed webfilter profiles. This is addressed in the Firewall Handbook.

The second method is the system-wide approach that locally (on the FortiGate Firewall) reassigns a URL to a different FortiGuard Category or even subcategory. This is where you can assign a specific URL to the FortiGuard Category that you want to you can also set the URL to one of the Custom Categories that you have created

The Web Rating Overrides option is available because different people will have different criteria for how they categorize websites. Even if the criteria is the same an organization may have reason to block the bulk of a category but need to be able to access specific URLs that are assigned to that category.

A hypothetical example could be that a website, example.com is categorized as being in the Sub-Category Pornography. The law offices of Barrister, Solicitor, and Lawyer do not want their employees looking at pornography at work so they have used the FortiGuard Webfilter to block access to sites that have been assigned to the Category “Pornography”. However, the owners of example.com are clients of the law office and they are aware that example.com is for artists that specialize in nudes and erotic images. In this case two approaches can be taken. The first is that the Web Rating Override function can be used to assign example.com to Nudity and Risque instead of Pornography for the purposes of matching the criteria that the law office goes by or the site can be assigned to a Custom Category that is not blocked because the site belongs to one of their clients and they always want to be able to access the site.

Another hypothetical example from the other side of the coin. A private school has decided that a company that specializes in the online selling of books that could be considered inappropriate for children because of their violent subject matter, should not be accessible to anyone in the school. The categorization by Fortinet of the site example2.com is General Interest - Business with the subcategory of Shopping and Auction, which is a category that is allowed at the school. In this case they school could reassign the site to the Category Adult Material which is a blocked category.

Local or Custom Categories

User-defined categories can be created to allow users to block groups of URLs on a per-profile basis. The categories defined here appear in the global URL category list when configuring a web filter profile. Users can rate URLs based on the local categories.

Users can create user-defined categories then specify the URLs that belong to the category. This allows users to block groups of web sites on a per profile basis. The ratings are included in the global URL list with associated categories and compared in the same way the URL block list is processed.



Local categories and local rating features consume a large amount of CPU resources; use these features as little as possible.

The local assignment of a category overrides the FortiGuard server ratings and appear in reports as “Local” Categories or “Custom” Categories depending on the context.

CLI commands

In the CLI, the term is local category.

To create a local category:

```
config webfilter ftgd-local-cat
  edit local_category_1
    set id 140
  end
```

To set a rating to a Local Category:

```
config webfilter ftgd-local-rating
edit <url_str>
    set rating {[<category_int>] [group_str] . . .]
    set status {enable | disable}
end
```

GUI commands

In the GUI, **Local Categories** appears on the **Edit Web Filter** profile page and **Custom Categories** on the **Web Rating Overrides** page, if your FortiGate is in proxy-based or flow-based, profile-based inspection. If your FortiGate is operating with flow-based inspection and the policy-based NGFW mode, then you will not see the **Edit Web Filter** profile page.

Both these features will be used to create local categories and to apply actions to them.

Creating a Local or Custom Category

1. Go to **Security Profiles > Web Rating Overrides**.
2. Select **Custom Categories** in the top menu bar.
3. In the new window, click on **Create New**.
4. Enter the name of the custom category.
5. Select **OK**.

Configuring Web Rating Overrides

Using the GUI

1. Go to **Security Profiles > Web Rating Overrides**.
2. Select **Create New**
3. Type in the **URL** field the URL of the Website that you wish to recategorize. Do not use wildcard expressions when typing in the URL.
4. Select the **Lookup Rating** button to verify the current categorization assigned to the URL.
5. Change the **Category** field to one of the more applicable options from the drop down menu, for example, one of the custom categories just created.
6. Change the **Sub-Category** field to a more narrowly defined option within the main category.
7. Select **OK**.



It is usually recommended that you choose a category that you know will be addressed in existing Web Filter profiles so that you will not need to engage in further configuration.

Applying an Action to a Local or Custom Category

1. Go to **Security Profiles > Web Filter**.
2. Expand the **Local Categories** in the list of FortiGuard categories.
3. Right-click on a category from the list and set the action to **Allow**, **Block**, **Monitor**, **Warning**, **Authenticate**, or **Disable**.
4. Select **Apply**.

You cannot apply an action to a local category when operating in flow-based NGFW policy-based mode.

Web filtering local and remote category status

The `status` option allows you to enable or disable FortiGuard web filtering category overrides for local and remote categories. When disabled, `ssl-exempt`, `webfilter`, and `proxy-address` cannot use the category. The status cannot be set to `disable` if it has been referenced.

Syntax

```
config webfilter ftgd-local-cat
  edit <name>
    set status {disable | disable}
    set id 140
  next
end
```

Local Category scenarios

Scenario 1: The configuration of the domain name overrides the configuration for the subdirectory.

Depending on the URL specified or other aspects of configuration, the configuration of a local or custom category may not take effect. Consider a scenario where you have defined:

- example.com – local rating as “category 1”, action set to Block
- example.com/subdirectory – local rating as “category 2”, action set to Monitor
- example.com/subdirectory/page.html – local rating as “category 3”, action set to Warning.

If a user browses to “example.com”, access will be blocked. If a user browses to example.com/subdirectory, access will also be blocked, even though that address was configured to be part of category2. The configuration of the domain name overrides the configuration for the subdirectory.

However, if you configure a specific HTML page differently than the domain name, then that configuration will apply. In this scenario, the user will see a Warning message but will be able to pass through to the page.

Scenario 2: User-defined local ratings and SNI matches

In this scenario, local categories are defined and sites are added to those categories.

- There is no behavioral difference if the hostname is sent from ClientHello SNI or from HTTP request-url.
- The SNI will be used as hostname for https certificate-inspection or ssl-exempt.
- If a valid SNI exist, then SNI will be used as the domain name for url rating instead of CN in the server certificate.
- For the local rating, “example.com” will match “test.example.com”, but will not match “another_example.com”.

Using Alternate Profiles

Allow Blocked Overrides or Web Overrides

The Administrative Override feature for Web Filtering is found by going to **Security Profiles > Web Filter** and then enabling **Allow users to override blocked categories**.

The Concept

When a Web filter profile is overridden, it does not necessarily remove all control and restrictions that were previously imposed by the Web Filter. The idea is to replace a restrictive filter with a different one. In practice, it

makes sense that this will likely be a profile that is less restrictive the original one but there is nothing that forces this. The degree to which the alternate profile is less restrictive is open. It can be as much as letting the user access everything on the Internet or as little as allowing only one additional website. The usual practice is to have as few alternate profiles as are needed to allow approved people to access what they need during periods when an exception to the normal rules is needed but still having enough control that the organizations web usage policies are not compromised.

You are not restricted to having only one alternative profile as an option to the existing profile. The new profile depends on the credentials or IP address making the connection. For example, John connecting through the "Standard" profile could get the "Allow_Streaming_Video" profile while George would get the "Allow_Social_Networking_Sites" profile.

The other thing to take into account is the time factor on these overrides. They are not indefinite. The longest that an override can be enabled is for 1 year less a minute. Often these overrides are set up for short periods of time for specific reasons such as a project. Having the time limitation means that the System Administrator does not have to remember to go back and turn the feature off after the project is finished.

Identity or Address

In either case, these override features -- for specified users, user groups or IP addresses -- allow sites blocked by Web Filtering profiles to be overridden for a specified length of time. The drawback of this method of override is that it takes more planning and preparation than the rating override method. The advantage is that once this has been set up, this method requires very little in the way of administrative overhead to maintain.

When planning to use the alternative profile approach keep in mind the following: In Boolean terms, one of the following "AND" conditions has to be met before overriding the Web Filter is possible.

Based on the IP address:

- The Web Filter profile must be specified as allowing overrides
- AND the user's computer is one of the IP addresses specified
- AND the time is within the expiration time frame.

While the conditions are fewer for this situation, there is less control over who has the ability to bypass the filtering configured for the site. All someone has to do is get on a computer that is allowed to override the Web Filter and they have access.

Based on user group:

- The Web Filter profile must be specified as allowing overrides
- AND the policy the traffic is going through must be identity based
- AND the user's credentials matches the identity credentials specified
- AND the time is within the expiration time frame.

This method is the one most likely to be used as it gives more control in that the user has to have the correct credential and is more versatile because the user can use the feature from any computer that uses the correct policy to get out on the Internet.

Settings

When using an alternate profile approach to Web Filter overrides, the following settings are used to determine authentication and outcome. Not every setting is used in both methods but enough of them are common to describe them collectively.

Apply to Group(s)

This is found in the Allow Blocked Overrides configuration. Individual users can not be selected. You can select one or more of the User Groups that are recognized by the FortiGate unit, whether they are local to the system or from a third part authentication device such as a AD server through FSSO.

Original Profile

This is found in the Administrative Override configuration. In the Allow Blocked Overrides setting the configuration is right inside the profile so there is no need to specify which profile is the original one, but the Administrative Override setup is done separately from the profiles themselves.

Assign to Profile or New Profile

Despite the difference in the name of the field, this is the same thing in both variations of the feature. You select from the drop down menu the alternate Web Filter Profile that you wish to set up for this override.

Scope or Scope Range

When setting up the override in the "Allow Blocked Overrides" variation, you are given a drop-down menu next to the field name Scope while in the Administrative Override configuration you are asked to select a radio button next to the same options. In both cases this is just a way of selecting which form of credentials will be required to approve the overriding of the existing Web Filter profile.

When the Web Filter Block Override message page appears it will display a field named "Scope:" and depending on the selection, it will show the type of credentials used to determine whether or not the override is allowed. The available options are:

- **User**
This means that the authentication for permission to override will be based on whether or not the user is using a specific user account.
- **User Group**
This means that the authentication for permission to override will be based on whether or not the user account supplied as a credential is a member of the specified User Group.
- **IP**
This means that the authentication for permission to override will be based on the IP address of the computer that was used to authenticate. This would be used with computers that have multiple users. Example: If Paul logs on to the computer, engages the override using his credentials and then logs off, if the scope was based on the IP address of the computer, anybody logging in with any account on that computer would now be using the alternate override Web Filter profile.

When entering an IP address in the Administrative Override version, only individual IP addresses are allowed.

Differences between IP and Identity based scope

- Using the IP scope does not require the use of an Identity based policy.
- When using the Administrative Override variation and IP scope, you may not see a warning message when you change from using the original Web Filter profile to using the alternate profile. There is no requirement for credentials from the user so, if allowed, the page will just come up in the browser.
- **Ask**
This option is available only in the "Allowed Blocked Overrides" variation and when used configures the message

page to ask which scope the user wished to use. Normally, when the page appears the scope options are greyed out and not editable, but by using the ask option the option is dark and the user can choose from the choice of:

- User
- User Group
- IP Address

• Switch Duration

The Administrative Override sets a specified time frame that is always used for that override. The available options are:

• Predefined

Using this setting will mean that what ever is set as the duration will be the length of time that the override will be in effect. If the duration variable is set to 15 minutes the length of the override will always be 15 minutes. The option will be visible in the Override message page but the setting will be greyed out.

• Ask

Using this setting will give the person the option of setting the duration to the override when it is engaged. The user can set the duration in terms of Day, Hours and or Minutes.

• Duration

Duration is one of the areas where the two variations take a different approach, on two aspects of the setting. As already indicated the "Administrative Override" only uses a static time frame there is no option for the user to select on the fly how long it will last. The other way in which the two variation differ is that the "Allow Blocked Overrides" starts the clock when the user logs in with his credentials. For example, if the duration is 1 hour and John initiates an override at 2:00 p.m. on January 1, at the end of that hour he will revert back to using the original profile but he can go back and re-authenticate and start the process over again. The Administrative override variation starts the clock from when the override was configured, which is why it shows an expiration date and time when you are configuring it.

This option, which is available when the Switch Mode is set to Predefined is the time in minutes that the override will last when engaged by the user.

When setting up a constant duration in the Web Based Interface, minutes is the only option for units of time. To set a longer time frame or to use the units of hours or days you can use the CLI.

```
config webfilter profile
edit <name of webfilter profile>
config override
set ovr-dur <###d##h##m>
end
```

When configuring the duration you don't have to set a value for a unit you are not using. If you are not using days or hours you can use:

```
set ovr-dur 30m
```

instead of:

```
set ovr-dur 0d0h30m
```

However, each of the units of time variable has their own maximum level:

```
###d cannot be more than 364
##h cannot be more than 23
##m cannot be more than 59
```

So the maximum length that the override duration can be set to is 364 days, 23 hours, and 59 minutes.



Using cookies to authenticate users in a Web Filter override

Cookies can be used to authenticate users when a web filter override is used. This feature is available in CLI only.

CLI syntax:

```
config webfilter cookie-ovrd
    set redir-host <name or IP>
    set redir-port <port>
end

config webfilter profile
    edit <name>
        config override
            set ovr-cookie [allow | deny]
            set ovr-scope [user | user-group | ip | ask]
            set profile-type [list | radius]
            set ovr-dur-mode [constant | ask]
            set ovr-dur <duration>
            set ovr-user-group <name>
            set profile <name>
        end
    end
end
```

Threat Feed Connectors

This feature introduces the ability to dynamically import external block lists in the form of a text file (containing a list of either addresses or domains), which reside on an HTTP server. These dynamic block lists are called 'Threat Feeds'. You can use threat feeds to deny access to a source or destination IP address in Web Filter and DNS Filter profiles, SSL inspection exemptions, and as a Source/Destination in proxy policies. The block list is stored as an external resource, which is dynamically imported to the FortiGate at a configured interval (or refresh-rate) in order to maintain an updated list.

In each profile, the administrator can configure multiple threat feeds.

Threat Feeds can be configured under **Security Fabric > Fabric Connectors** by creating new **Threat Feeds**.

The **External Resources** edit page provides the following fields:

- **Type**
 - **FortiGuard Category** - The resource **Name** will appear as a "Remote Category" in Web Filter profiles and SSL inspection exemptions.
 - **IP Address** - The resource **Name** will appear as an "External Domain Block List" in DNS Filter profiles and as a "Source/Destination" in proxy policies.
 - **Domain Name** - The resource **Name** will appear as an "External Domain Block List" in DNS Filter profiles.
- **URI of external resource** - The link to an external resource file. The file should be a plain text file with one domain each line and supports simple wildcard.
- **Refresh Rate** - The time interval to refresh external resource (1 - 43200 minutes).
- The size of the file can be 10 MB, or 128,000 lines of text, whichever is most restrictive.

The domain resource is a text file which contains a domain name for each line and supports simple wildcard. For example:

```
mail.*.or.th
*-special.de.vu
http://www.*de.vu
610-pawn.com
aaliyah-hq-gallery.de.vu
abcgolocal.com
```

The address resource is a text file which contains an IP/IP range for each line (note that only IPv4 is supported in DNS profiles, so IPv6 addresses will be ignored). For example:

```
1.1.1.1
10.0.0.70
2.1.1.1
100.0.0.1-100.0.0.100
10.0.0.99-10.0.0.201
1.2.2.2/24
```

Syntax

```
config system external-resource
  edit <name>
    set type {category | address | domain}
    set category <value>
    set comments [comments]
    set resource <resource-url>
    set refresh-rate <minutes>
    set last-update <datetime>
  next
end
```

You can also configure one or more domain list threat feeds under `config dnsfilter profile`. See ["DNS filter" on page 2482](#) for more information.

Web Profile Overrides

This feature allows administrators to grant temporary access to sites that are otherwise blocked by a web filter profile. The temporary access can be granted to a user, user group, or source IP address. The time limit can be set in days, hours, or minutes. The default is 15 minutes.

Temporary access can also be granted to a user, user group, or source IP address by enabling **Allow users to override blocked categories** in a Web Filter security profile and applying that profile to the appropriate policy. In this scenario, the user will have to authenticate to gain access.

When Web Profile Overrides is in effect, a blocked access page or replacement message will not appear and authentication will not be required.



If your FortiGate is operating in flow-based inspection and policy-based NGFW mode, then you cannot create a web profile override.

Creating a Web Profile Override

Before creating a Web Profile Override, you will have to configure a user or user group if not granting temporary access to a Source IP. You will also have to configure the Web Filter security profile to be applied to the override.

1. Go to **Security Profiles > Web Profile Overrides**.
2. Select **Create New**.

3. Select:
 - The **Scope Range: User** and **User Group** should be previously configured under **User & Device**
 - The **Original Profile** that applied to the scope range.
 - The **New Profile** to apply for the override
 - The time when the override **Expires**; default is 15 minutes

SafeSearch

SafeSearch is a feature of popular search sites that prevents explicit web sites and images from appearing in search results. Although SafeSearch is a useful tool, especially in educational environments, the resourceful user may be able to simply turn it off. Enabling SafeSearch for the supported search sites enforces its use by rewriting the search URL to include the code to indicate the use of the SafeSearch feature. For example, on a Google search it would mean adding the string "&safe=active" to the URL in the search.

The search sites supported are:

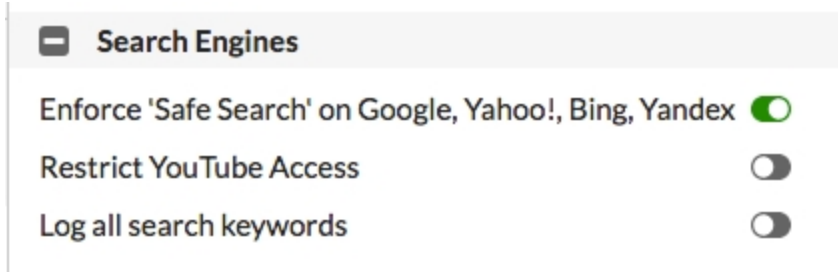
- Google
- Yahoo
- Bing
- Yandex



You can only enable SafeSearch with proxy-based inspection mode.

Enabling SafeSearch - GUI

1. Go to **Security Profiles > Web Filter** and edit or create a policy.
2. Expand **Search Engines**



Enabling SafeSearch - CLI

```
config webfilter profile
  edit default
    config web
      set safe-search <url>
    end
  end
end
```

This enforces the use of SafeSearch in traffic controlled by the firewall policies using the web filter you configure.

Search Keywords

There is also the capability to log the search keywords used in the search engines.

YouTube Education Filter

YouTube for Schools was a way to access educational videos from inside a school network. This YouTube feature gave schools the ability to access a broad set of educational videos on YouTube EDU and to select the specific videos that are accessible from within the school network.

Google decided to stop supporting YouTube for Schools (YTfS) as of July 1, 2016. Consequently, the current YouTube safe search does not work anymore.

Google provides an article entitled "[Restrict YouTube content on your network or managed devices](#)" on its support site. At this time, Google offers two options to restrict inappropriate content: DNS and HTTP header.

To restrict YouTube access, go to **Security Profiles > Web Filter**, scroll to **Search Engines** and enable **Restrict YouTube Access**. You can select either **Strict** or **Moderate** level of restriction. Your FortiGate must be in proxy mode.

YouTube Channel Filtering

This Web Filtering feature lets you block or allow matched YouTube channels using one of the following identifiers:

- **<channel-id>**
- **www.youtube.com/channel/<channel-id>**
- **www.youtube.com/user/<user-id>**
matches channel-id from <meta itemprop="channelId" content="UCGzuiiLdQZu9wxDNJHO_JnA">

- **www.youtube.com/watch?v=<string>**
matches channel-id from <meta itemprop="channelId" content="UCGzuilLdQZu9wxDNJHO_JnA">

Syntax

Note that `config youtube-channel-filter` is only available when `youtube-channel-status` is set to either `blacklist` or `whitelist`. Also note that, when defining `channel-id`, both the full URL or just the Channel ID suffix are acceptable, as shown below:

```
config webfilter profile
  edit <name>
    set youtube-channel-status {disable | blacklist | whitelist}
    config youtube-channel-filter
      edit <id>
        set channel-id <url>
      next
      edit <id>
        set channel-id <channel-id>
      next
    end
  end
end
```

Static URL filter

You can allow or block access to specific URLs by adding them to the **Static URL Filter** list. The filter allows you to block, allow, or monitor URLs by using patterns containing text, regular expressions, or wildcard characters. The FortiGate unit allows or blocks web pages matching any specified URLs or patterns and displays a replacement message instead.



URL blocking does not block access to other services that users can access with a web browser. For example, URL blocking does not block access to `ftp://ftp.example.com`. Instead, use firewall policies to deny ftp connections.

When adding a URL to the URL filter list, follow these rules:

- Type a top-level URL or IP address to control access to all pages on a web site. For example, `www.example.com` or `192.168.144.155` controls access to all pages at this web site.
- Enter a top-level URL followed by the path and file name to control access to a single page on a web site. For example, `www.example.com/news.html` or `192.168.144.155/news.html` controls access to the news page on this web site.
- To control access to all pages with a URL that ends with `example.com`, add `example.com` to the filter list. For example, adding `example.com` controls access to `www.example.com`, `mail.example.com`, `www.finance.example.com`, and so on.
- Control access to all URLs that match patterns using text and regular expressions (or wildcard characters). For example, `example.*` matches `example.com`, `example.org`, `example.net` and so on.



URLs with an action set to exempt or monitor are not scanned for viruses. If users on the network download files through the FortiGate unit from a trusted web site, add the URL of this web site to the URL filter list with an action to pass it so the FortiGate unit does not virus scan files downloaded from this URL.

URL formats

How URL formats are detected when using HTTPS

Filter HTTPS traffic by entering a top level domain name, for example, `www.example.com` if:

- your unit does not support SSL content scanning and inspection
- you have selected the **URL filtering** option in web content profile for **HTTPS content filtering mode** under **Protocol Recognition**.

HTTPS URL filtering of encrypted sessions works by extracting the CN from the server certificate during the SSL negotiation. Since the CN only contains the domain name of the site being accessed, web filtering of encrypted HTTPS sessions can only filter by domain names.

If your unit supports SSL content scanning and inspection and if you have selected Deep Scan, you can filter HTTPS traffic in the same way as HTTP traffic.

How URL formats are detected when using HTTP

URLs with an action set to Exempt are not scanned for viruses. If users on the network download files through the unit from trusted web site, add the URL of this web site to the URL filter list with an action set to exempt so the unit does not virus scan files downloaded from this URL.

- Type a top-level URL or IP address to control access to all pages on a web site. For example, `www.example.com` or `192.168.144.155` controls access to all pages at this web site.
- Enter a top-level URL followed by the path and filename to control access to a single page on a web site. For example, `www.example.com/news.html` or `192.168.144.155/news.html` controls the news page on this web site.
- To control access to all pages with a URL that ends with `example.com`, add `example.com` to the filter list. For example, adding `example.com` controls access to `www.example.com`, `mail.example.com`, `www.finance.example.com`, and so on.
- Control access to all URLs that match patterns created using text and regular expressions (or wildcard characters). For example, `example.*` matches `example.com`, `example.org`, `example.net` and so on.
- Fortinet URL filtering supports standard regular expressions.



If virtual domains are enabled on the unit, web filtering features are configured globally. To access these features, select **Global Configuration** on the main menu.

URL filter actions

You can select one of four actions for how traffic will be treated as it attempts to reach a site in the list.

Block

Attempts to access any URLs matching the URL pattern are denied. The user will be presented with a replacement message.

Allow

Any attempt to access a URL that matches a URL pattern with an allow action is permitted. The traffic is passed to the remaining antivirus proxy operations, including FortiGuard Web Filter, web content filter, web script filters,

and antivirus scanning.

Allow is the default action. If a URL does not appear in the URL list, it is permitted.

Monitor

Traffic to, and reply traffic from, sites matching a URL pattern with **Monitor** action applied will be allowed through in the same way as the **Allow** action. The difference with the **Monitor** action is that a log message will be generated each time a matching traffic session is established. The requests will also be subject to all other Security Profiles inspections that would normally be applied to the traffic.

Exempt

Exempt allows trusted traffic to bypass the antivirus and DLP proxy operations by default, but it functions slightly differently. In general, if you're not certain that you need to use the **Exempt** action, use **Monitor**.



Using the static URL filter to exempt scanning also prevents SSL inspection.

HTTP 1.1 connections are persistent unless declared otherwise. This means the connections will remain in place until closed or the connection times out. When a client loads a web page, the client opens a connection to the web server. If the client follows a link to another page on the same site before the connection times out, the same connection is used to request and receive the page data.

When you add a URL pattern to a URL filter list and apply the **Exempt** action, traffic sent to and replies traffic from sites matching the URL pattern will bypass all antivirus proxy operations. The connection itself inherits the exemption. This means that all subsequent reuse of the existing connection will also bypass all antivirus proxy operations. When the connection times out, the exemption is cancelled.

For example, consider a URL filter list that includes `example.com/files` configured with the **Exempt** action. A user opens a web browser and downloads a file from the URL `example.com/sample.zip`. This URL does not match the URL pattern so it is scanned for viruses. The user then downloads `example.com/files/beautiful.exe` and since this URL does match the pattern, the connection itself inherits the exempt action. The user then downloads `example.com/virus.zip`. Although this URL does not match the exempt URL pattern, a previously visited URL did, and since the connection inherited the exempt action and was re-used to download a file, the file is not scanned.

If the user next goes to an entirely different server, like `example.org/photos`, the connection to the current server cannot be reused. A new connection to `example.org` is established. This connection is not exempt. Unless the user goes back to `example.com` before the connection to that server times out, the server will close the connection. If the user returns after the connection is closed, a new connection to `example.com` is created and it is not exempt until the user visits a URL that matches the URL pattern.

Web servers typically have short time-out periods. A browser will download multiple components of a web page as quickly as possible by opening multiple connections. A web page that includes three photos will load more quickly if the browser opens four connections to the server and downloads the page and the three photos at the same time. A short time-out period on connections will close the connections faster, allowing the server to avoid unnecessarily allocating resources for a long period. The HTTP session time-out is set by the server and will vary with the server software, version, and configuration.

Using the **Exempt** action can have unintended consequences in certain circumstances. You have a web site at `example.com` and since you control the site, you trust the contents and configure `example.com` as exempt. But

example.com is hosted on a shared server with a dozen other different sites, each with a unique domain name. Because of the shared hosting, they also share the same IP address. If you visit example.com, your connection to your site becomes exempt from any antivirus proxy operations. Visits to any of the 12 other sites on the same server will reuse the same connection and the data you receive is exempt from being scanned.

Use of the **Exempt** action is not suitable for configuration in which connections through the FortiGate unit use an external proxy. For example, you use proxy.example.net for all outgoing web access. Also, as in the first example, URL filter list that includes a URL pattern of `example.com/files` configured with the **Exempt** action. Users are protected by the antivirus protection of the FortiGate unit until a user visits a URL that matches the `example.com/files` URL pattern. The pattern is configured with the **Exempt** action so the connection to the server inherits the exemption. With a proxy however, the connection is from the user to the proxy. Therefore, the user is entirely unprotected until the connection times out, no matter what site he visits.

Ensure you are aware of the network topology involving any URLs to which you apply the **Exempt** action.

Status

The Web Site Filter has the option to either enable or disable individual web sites in the list. This allows for the temporary removal of the actions against a site so that it can be later reengaged without having to rewrite the configuration.

Configuring a URL filter

Consult the [Maximum Values Table](#) on the [Fortinet Document Library](#) site for up-to-date information on the number of URL filter entries allowed for your FortiGate.



You can only set a Static URL Filter with proxy-based inspection mode and flow-based inspection mode in profile-based NGFW mode.

For this example, the URL `www.example*.com` will be used. You configure the list by adding one or more URLs to it.

To add a URL to a URL filter

1. Go to **Security Profiles > Web Filter**.
2. Create a new web filter or select a one to edit.
3. Expand **Static URL Filter**, enable **URL Filter**, and select **Create**.
4. Enter the URL, without the "http", for example: `www.example*.com`.
5. Select a **Type**: **Simple**, **Reg. Expression**, or **Wildcard**. In this example, select **Wildcard**.
6. Select the **Action** to take against matching URLs: **Exempt**, **Block**, **Allow**, or **Monitor**.
7. Confirm that **Status** is enabled.
8. Select **OK**.

'Simple' Filter type

If you select the **Simple** filter type for a URL filter, the syntax is performing an exact match. Note, however, that the domain and path are separate entities in HTTP despite the fact that a user types them as a single entity and, in the case of 'simple', the rules for each part (domain and path) are different.

The 'domain' part

For the domain part, the goal of the 'simple' format is to make it easy to block a domain and all its subdomains, such that the admin only has to type "address.xy" to block "address.xy", "www.address.xy", "talk.address.xy", etc. but *not* block "youraddress.xy" or "www.youraddress.xy" which are different domains from "address.xy".

Also, the actual domain does not include http:// or https:// so this should *not* be entered or the URL filter will try to match a domain starting with http. For this reason, when you enter http:// in the URL filter via the GUI, it is automatically removed.



A trailing '/' with the domain is not needed. The GUI URL filter will automatically trim this, but when using the API to provide the per-user BWL it will not!

Please take this into account. Better not to use it as it might give unexpected results.

The 'path' part

For the path part, an exact match takes place. For example:

www.address.xy/news

blocks anything that starts with that exact path. So this matches:

www.address.xy/newsies
www.address.xy/newsforyou
www.address.xy/news/co
etc.

Also:

www.address.xy/new

likewise blocks the same as above but includes:

/newt
/newp
etc.

which is a much broader filter, matching:

www.address.xy/newstand/co
www.address.xy/news/co
etc.

In other words, the more you specify of the path, the more strictly it will match.



Here as well a trailing '/' with the URL path is not needed, the GUI URL filter will automatically trim this, but when using the API to provide the per-user BWL it will not!

Please take this into account. Better not to use it as it might give unexpected results.

Referrer URL

A new variable has been added to the Static URL Filter: `referrer-host`. If a referrer is specified, the hostname in the referrer field of the HTTP require will be compared for any entry that contains the matching URL.

If the referrer matches, then the specified action will be performed by proxy.

Configuring in the GUI

The configuration can be done in the GUI but only if advance web filtering features have been enabled by entering the following commands in the CLI:

```
config system global
    set gui-webfilter-advanced enable
end
```

After this command is used, a new column will be created in **Security Profiles > Web Filter** to set the referrer.

Configuring in the CLI

When specifying the URL filter, it needs to be identified by its ID. The URLs are listed under each entry.

To find the ID number:

```
config webfilter urlfilter
    edit ?
```

A list of the current URL filters will be listed with their ID numbers in the left column.

The syntax in the CLI for configuring an entry is:

```
config webfilter urlfilter
    edit <ID>
        config entries
            edit 1
                set url <url>
                set referrer-host <url>
                set type {simple | regex | wildcard}
                set action {block | allow | monitor | exempt}
                set status {enable | disable}
            end
        end
    end
```

Web content filter

You can control web content by blocking access to web pages containing specific words or patterns. This helps to prevent access to pages with questionable material. You can also add words, phrases, patterns, wild cards and Perl regular expressions to match content on web pages. You can add multiple web content filter lists and then select the best web content filter list for each web filter profile.

Enabling web content filtering involves three separate parts of the FortiGate configuration.

- The security policy allows certain network traffic based on the sender, receiver, interface, traffic type, and time of day.
- The web filter profile specifies what sort of web filtering is applied.
- The web content filter list contains blocked and exempt patterns.

The web content filter feature scans the content of every web page that is accepted by a security policy. The system administrator can specify banned words and phrases and attach a numerical value, or score, to the importance of those words and phrases. When the web content filter scan detects banned content, it adds the

scores of banned words and phrases in the page. If the sum is higher than a threshold set in the web filter profile, the FortiGate unit blocks the page.

General configuration steps

Follow the configuration procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. Create a web content filter list.
2. Add patterns of words, phrases, wildcards, and regular expressions that match the content to be blocked or exempted.
3. You can add the patterns in any order to the list. You need to add at least one pattern that blocks content.
4. In a web filter profile, enable the web content filter and select a web content filter list from the options list.

To complete the configuration, you need to select a security policy or create a new one. Then, in the security policy, enable **Webfilter** and select the appropriate web filter profile from the list.

Creating a web filter content list

You can create multiple content lists and then select the best one for each web filter profile. Creating your own web content lists can be accomplished only using the CLI.

This example shows how to create a web content list called inappropriate language, with two entries, offensive and rude.

To create a web filter content list

```
config webfilter content
  edit 3
    set name "inappropriate language"
    config entries
      edit offensive
        set action block
        set lang western
        set pattern-type wildcard
        set score 15
        set status enable
      next
      edit rude
        set action block
        set lang western
        set pattern-type wildcard
        set score 5
        set status enable
      end
    end
  end
end
```

Configuring a web content filter list

Once you have created the web filter content list, you need to add web content patterns to it. There are two types of patterns: **Wildcard** and **Regular Expression**.

You use the **Wildcard** setting to block or exempt one word or text strings of up to 80 characters. You can also use the wildcard symbols, such as "*" or "?", to represent one or more characters. For example, as a wildcard

expression, `forti*.com` will match `fortinet.com` and `forticare.com`. The `"*"` represents any kind of character appearing any number of times.

You use the **Regular Expression** setting to block or exempt patterns of Perl expressions, which use some of the same symbols as wildcard expressions, but for different purposes. The `"*"` represents the character before the symbol. For example, `forti*.com` will match `fortiii.com` but not `fortinet.com` or `fortiice.com`. The symbol `"i"` represents "i" in this case, appearing any number of times. RP: Add a regex example.

The maximum number of web content patterns in a list is 5000.

How content is evaluated

Every time the web content filter detects banned content on a web page, it adds the score for that content to the sum of scores for that web page. You set this score when you create a new pattern to block the content. The score can be any number from zero to 99999. Higher scores indicate more offensive content. When the sum of scores equals or exceeds the threshold score, the web page is blocked. The default score for web content filter is 10 and the default threshold is 10. This means that by default a web page is blocked by a single match. Blocked pages are replaced with a message indicating that the page is not accessible according to the Internet usage policy.

Banned words or phrases are evaluated according to the following rules:

- The score for each word or phrase is counted only once, even if that word or phrase appears many times in the web page.
- The score for any word in a phrase without quotation marks is counted.
- The score for a phrase in quotation marks is counted only if it appears exactly as written.

The following table describes how these rules are applied to the contents of a web page. Consider the following, a web page that contains only this sentence: "The score for each word or phrase is counted only once, even if that word or phrase appears many times in the web page."

Banned pattern rules

Banned pattern	Assigned score	Score added to the sum for the entire page	Threshold score	Comment
word	20	20	20	Appears twice but only counted once. Web page is blocked.
word phrase	20	40	20	Each word appears twice but only counted once giving a total score of 40. Web page is blocked
word sentence	20	20	20	"word" appears twice, "sentence" does not appear, but since any word in a phrase without quotation marks is counted, the score for this pattern is 20. Web page is blocked.
"word sentence"	20	0	20	"This phrase does not appear exactly as written. Web page is allowed.
"word or phrase"	20	20	20	This phrase appears twice but is counted only once. Web page is blocked.

Enabling the web content filter and setting the content threshold

When you enable the web content filter, the web filter will block any web pages when the sum of scores for banned content on that page exceeds the content block threshold. The threshold will be disregarded for any exemptions within the web filter list.

Web filtering example

Web filtering is particularly important for protecting school-aged children. There are legal issues associated with improper web filtering as well as a moral responsibility to keep children from viewing inappropriate material. The key is to design a web filtering system in such a way that students and staff do not fall under the same web filter profile in the FortiGate configuration. This is important because the staff may need to access websites that are off limits to the students.

School district

The background for this scenario is a school district with more than 2,300 students and 500 faculty and staff in a preschool, three elementary schools, a middle school, a high school, and a continuing education center. Each elementary school has a computer lab and the high school has three computer labs with connections to the Internet. Such easy access to the Internet ensures that every student touches a computer every day.

With such a diverse group of Internet users, it was not possible for the school district to set different Internet access levels. This meant that faculty and staff were unable to view websites that the school district had blocked. Another issue was the students' use of proxy sites to circumvent the previous web filtering system. A proxy server acts as a go-between for users seeking to view web pages from another server. If the proxy server has not been blocked by the school district, the students can access the blocked website.

When determining what websites are appropriate for each school, the district examined a number of factors, such as community standards and different needs of each school based on the age of the students.

The district decided to configure the FortiGate web filtering options to block content of an inappropriate nature and to allow each individual school to modify the options to suit the age of the students. This way, each individual school was able to add or remove blocked sites almost immediately and have greater control over their students' Internet usage.

In this simplified example of the scenario, the district wants to block any websites with the word **example** on them, as well as the website www.example.com. The first task is to create web content filter lists for the students and the teachers.



Web Filter in flow mode does not support Safe Search. The examples below use Web Filter in proxy mode.

Create a Web Filter profile for the students

1. Go to **Security Profiles > Web Filter**.
2. Select the **Create New** icon.
3. Enter the name "Students" in the name field.
4. Enable FortiGuard Categories.
 - a. Set the following categories to **Block**:

- Potentially Liable
- Adult/Mature Content
- Security Risk

URL Content

6. Go to **Search Engines** and expand the section if necessary. Enable **Enforce 'Safe Search' on Google, Yahoo!, Bing, Yandex**
7. In the **Static URL Filter** section, enable **URL Filter**.
 - a. Select **Create New**.
 - i. In the **URL** field, enter ***example*.***
 - ii. For the **Type** field, select **Wildcard**
 - iii. For the **Action** field, select **Block**
 - iv. For the **Status** field, check **enable**
 - v. Select **OK**

Web Content Filter

8. In the **Static URL Filter** section, enable **Web Content Filter**.
 - a. In the **Web Content Filter** widget, select **Create New**.
 - i. For the **Pattern Type** field, select **Reg. Expression**
 - ii. In the **Pattern** field, enter "example"
 - iii. For the **Language** field, choose Western
 - iv. For the **Action** field, select "Block"
 - v. For the **Status** field, check Enable.
 - vi. Select **OK**
9. Enable **Rate URLs by Domain and IP Address**
10. Disable **Allow websites when a rating error occurs**.
11. Check **Rate Images by URL (Blocked images will be replaced with blanks)**
12. Select **Apply**

Create a Web Filter for the teachers

It might be more efficient if the Teacher Web Content List included the same blocked content as the student list. From time to time a teacher might have to view a blocked page. It would then be a matter of changing the **Action** from **Block** to **Allow** as the situation required. The following filter is how it could be set up for the teachers to allow them to see the "example" content if needed while keeping the blocking inappropriate material condition.

1. Go to **Security Profiles > Web Filter**.
2. Select the **Create New** icon.
3. Enter the name "Teachers" in the name field.
4. Enable FortiGuard Categories.
 - a. Set the following categories to **Block**:
 - Potentially Liable
 - Adult/Mature Content
 - Security Risk

URL Content

6. Go to **Search Engines** and expand the section if necessary. Enable **Search Engine Safe Search on Google, Yahoo!, Bing, Yandex**.
7. In the Static URL Filter section, check **Enable URL Filter**.
 - a. Select **Create New**.
 - i. In the **URL** field, enter *example*.*
 - ii. For the **Type** field, select **Wildcard**
 - iii. For the **Action** field, select Block
 - iv. For the **Status** field, check enable
 - v. Select **OK**

Web Content Filter

8. In the Static URL Filter section, check Enable Web Content Filter.
 - a. In the Web Content Filter widget, select **Create New**.
 - b. Enter the name "Teachers" in the name field.
 - i. For the **Pattern Type** field, select **Reg. Expression**
 - ii. In the **Pattern** field, enter "example"
 - iii. For the **Language** field, choose Western
 - iv. For the **Action** field, select Exempt
 - v. For the **Status** field, check Enable.
 - vi. Select **OK**
9. Check **Rate URLs by Domain and IP Address**
10. Check **Rate Images by URL (Blocked images will be replaced with blanks)**
11. Select **OK**

To create a security policy for the students

1. Go to **Policy & Objects > IPv4 Policy**.
2. Select the policy being used to manage student traffic.
3. Enable **Web Filter**.
4. Select **Students** from the web filter drop-down list.
5. Select **OK**.

To create a security policy for Teachers

1. Go to **Policy & Objects > IPv4 Policy**.
2. Select the policy being used to manage teacher traffic.
3. Enable **Web Filter**.
4. Select **Teachers** from the web filter drop-down list.
5. Select **OK**.
6. Make sure that the student policy is in the sequence before the teachers' policy.

Advanced web filter configurations

Allow websites when a rating error occurs

Enable this setting to allow access to web pages that return a rating error from the FortiGuard Web Filter service.

If your FortiGate unit cannot contact the FortiGuard service temporarily, this setting determines the type of access the FortiGate unit allows until contact is re-established. If enabled, users will have full unfiltered access to all web sites. If disabled, users will not be allowed access to any web sites.

ActiveX filter

Enable to filter ActiveX scripts from web traffic. Web sites using ActiveX may not function properly with this filter enabled.

Block HTTP redirects by rating

Enable to block HTTP redirects.

Many web sites use HTTP redirects legitimately but in some cases, redirects may be designed specifically to circumvent web filtering, as the initial web page could have a different rating than the destination web page of the redirect.

This option is not supported for HTTPS.

Block Invalid URLs

Select to block web sites when their SSL certificate CN field does not contain a valid domain name.

FortiGate units always validate the CN field, regardless of whether this option is enabled. However, if this option is not selected, the following behavior occurs:

- If the request is made directly to the web server, rather than a web server proxy, the FortiGate unit queries for FortiGuard Web Filtering category or class ratings using the IP address only, not the domain name.
- If the request is to a web server proxy, the real IP address of the web server is not known. Therefore, rating queries by either or both the IP address and the domain name is not reliable. In this case, the FortiGate unit does not perform FortiGuard Web Filtering.



Enabling the Web Filter profile to block a particular category and enabling the Application Control profile will not result in blocking the URL. This occurs because proxy and flow-based profiles cannot operate together.

To ensure replacement messages show up for blocked URLs, switch the Web Filter to flow-based inspection.

Cookie filter

Enable to filter cookies from web traffic. Web sites using cookies may not function properly with this enabled.

Provide Details for Blocked HTTP 4xx and 5xx Errors

Enable to have the FortiGate unit display its own replacement message for 400 and 500-series HTTP errors. If the server error is allowed through, malicious or objectionable sites can use these common error pages to

circumvent web filtering.

HTTP POST action

Select the action to take with HTTP POST traffic. HTTP POST is the command used by your browser when you send information, such as a form you have filled-out or a file you are uploading, to a web server.

The available actions include:

Comfort

Use client comforting to slowly send data to the web server as the FortiGate unit scans the file. Use this option to prevent a server time-out when scanning or other filtering is enabled for outgoing traffic.

The client comforting settings used are those defined in the Proxy Options profile selected in the security policy.

Block

Block the HTTP POST command. This will limit users from sending information and files to web sites.

When the post request is blocked, the FortiGate unit sends the http-post-block replacement message to the web browser attempting to use the command.

Java applet filter

Enable to filter java applets from web traffic. Web sites using java applets may not function properly with this filter enabled.

Rate Images by URL

Enable to have the FortiGate retrieve ratings for individual images in addition to web sites. Images in a blocked category are not displayed even if they are part of a site in an allowed category.

Blocked images are replaced on the originating web pages with blank place-holders. Rated image file types include GIF, JPEG, PNG, BMP, and TIFF.

Rate URLs by Domain and IP Address

Enable to have the FortiGate unit request the rating of the site by URL and IP address separately, providing additional security against attempts to bypass the FortiGuard Web Filter.

If the rating determined by the domain name and the rating determined by the IP address defer the Action that is enforce will be determined by a weighting assigned to the different categories. The higher weighted category will take precedence in determining the action. This will have the side effect that sometimes the Action will be determined by the classification based on the domain name and other times it will be determined by the classification that is based on the IP address.



FortiGuard Web Filter ratings for IP addresses are not updated as quickly as ratings for URLs. This can sometimes cause the FortiGate unit to allow access to sites that should be blocked, or to block sites that should be allowed.

An example of how this would work would be if a URL's rating based on the domain name indicated that it belonged in the category Lingerie and Swimsuit, which is allowed but the category assigned to the IP address was

Pornography which has an action of Block, because the Pornography category has a higher weight the effective action is Block.

Web resume download block

Enable to prevent the resumption of a file download where it was previously interrupted. With this filter enabled, any attempt to restart an aborted download will download the file from the beginning rather than resuming from where it left off.

This prevents the unintentional download of viruses hidden in fragmented files.

Note that some types of files, such as PDF, fragment files to increase download speed and enabling this option can cause download interruptions. Enabling this option may also break certain applications that use the Range Header in the HTTP protocol, such as YUM, a Linux update manager.

Restrict Google account usage to specific domains

This feature allow the blocking of access to some Google accounts and services while allowing access to accounts that are included in the domains specified in the exception list.

Block non-English character URLs

The FortiGate will not successfully block non-English character URLs if they are added to the URL filter. In order to block access to URLs with non-English characters, the characters must be translated into their international characters.

Browse to the non-English character URL (for example, <http://www.fortinet.com/pages/ที่'นี้'-ไม่มี'เคอร์'รูประหารให้'ใคร'ตก/338419686287505?ref=stream>).

On the FortiGate, use the URL shown in the FortiGate GUI and add it the list of blocked URLs in your URL filter (for example,

<http://www.fortinet.com/pages/%E0%B8%97%E0%B8%B5%E0%B9%88%E0%B8%99%E0%B8%B5%E0%B9%88-%E0%B9%84%E0%B8%A1%E0%B9%88%E0%B8%A1%E0%B8%B5%E0%B9%80%E0%B8%A8%E0%B8%A9%E0%B8%A3%E0%B8%B1%E0%B8%90%E0%B8%9B%E0%B8%A3%E0%B8%B0%E0%B8%AB%E0%B8%B2%E0%B8%A3%E0%B9%83%E0%B8%AB%E0%B9%89%E0%B9%83%E0%B8%84%E0%B8%A3%E0%B9%81%E0%B8%94%E0%B8%81/338419686287505?ref=stream>).

Once added, further browsing to the URL will result in a blocked page.

CLI Syntax

```
config webfilter urlfilter
edit 1
set name "block_international_character_urls"
config entries
edit 1
set url "www.fortinet.com/pages/2.710850E-3120%B8%E0%B8%B53.231533E-3170%B9%E0%B8%E0%B8%B53.231533E-3170%B9%88-3.230415E-3170%B9%E0%B80X0.000000063CD94P-102211.482197E-3230%B9%E0%B80X0.0007FBFFFFCFP-102210.000000E+000%B8%B51.828043E-3210%B9%E0%B80X0P+081.828043E-3210%B80X0P+092.710850E-3120%B80X0.0000000407ED2P-102233.236834E-3170%B8%B19.036536E-3130%B8%E0%B8%9B4.247222E-3140%B80X0P+039.036683E-3130%B8%B02.121996E-3130%B80X0.0000000000008P-1022B2.710850E-3120%B8%B21.482197E-3230%B80X0P+030.000000E+000%B9%E0%B80X0P+0B2.710850E-3120%B9%E0%B9%E0%B8%E0%B80X0.0000000408355P-102232.023693E-3200%B9%E0%B8%E0%B8%81/338419686287505?ref=stream"
```

```

        set action block
      next
    end
  next
end

config webfilter urlfilter
  edit 2
    set name "block_international_character_urls"
  next
end

config webfilter profile
  edit "block_international_character_urls"
  next
end

config firewall policy
  edit 3
    set uuid cf80d386-7bcf-51e5-6e87-db207e3f0fa8
    set srcintf "port1"
    set dstintf "port2"
    set srcaddr "all"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ALL"
    set utm-status enable
    set logtraffic all
    set webfilter-profile "block_international_character_urls"
    set profile-protocol-options "default"
    set ssl-ssh-profile "certificate-inspection"
    set nat enable
  next
end

```

Websense web filtering through WISP

WISP is a Websense protocol that allows for URLs to be extracted by a firewall and submitted to Websense systems for rating and approval checking.

This feature provides a solution for customers who have large, existing, deployed implementations of Websense security products to replace their legacy firewalls with a FortiGate family, such that they are not forced to make a change to their web filtering infrastructure at the same time.

When WISP is enabled, the FortiGate will maintain a pool of TCP connections to the WISP server. The TCP connections will be used to forward HTTP request information and log information to the WISP server and receive policy decisions.

Configuring the WISP server

In order to use WebSense's web filtering service, a WISP server per VDOM must be defined and enabled first.

```

config web-proxy wisp
  edit {name}
    # Configure Wireless Internet service provider (WISP) servers.

```

```
set name {string}    Server name. size[35]
set comment {string}  Comment. size[255]
set outgoing-ip {ipv4 address any}    WISP outgoing IP address.
set server-ip {ipv4 address any}    WISP server IP address.
set server-port {integer}    WISP server port (1 - 65535, default = 15868). range[1-65535]
set max-connections {integer}    Maximum number of web proxy WISP connections (4 - 4096, default =
64). range[4-4096]
set timeout {integer}    Period of time before WISP requests time out (1 - 15 sec, default = 5). range
[1-15]
next
end
```

Example configuration

```
config web-proxy wisp
edit 0
set outgoing-ip 0.0.0.0
set server-ip 0.0.0.0
set server-port 15868
set max-connections 64
set timeout 5
next
end
```

After configuring the WISP server, enable WISP in the web filter profile.

```
config webfilter profile
edit "wisp_only"
set wisp enable
set wisp-servers 0
next
end
```

Now you can apply the web filter profile to a firewall policy.

If you configure more than one WISP server, the load balance option can also be configured.

```
config webfilter profile
edit "wisp_only"
set wisp-algorithm {primary-secondary | round-robin | auto-learning}
next
end
```

The options for the wisp-algorithm are:

- primary-secondary: select the first healthy server in order
- round-robin: select the next healthy server
- auto-learning select the lightest loading healthy server

DNS filter

You can configure DNS web filtering to allow, block, or monitor access to web content according to FortiGuard categories. When DNS web filtering is enabled, your FortiGate must use the FortiGuard DNS service for DNS lookups. DNS lookup requests sent to the FortiGuard DNS service return with an IP address and a domain rating that includes the FortiGuard category of the web page.

If that FortiGuard category is set to **block**, the result of the DNS lookup is not returned to the requester. If the category is set to **redirect**, then the address returned to the requester points at a FortiGuard redirect page.

You can also allow or monitor access based on FortiGuard category.

Blocking DNS requests to known botnet command & control addresses

FortiGuard maintains a database containing a list of known botnet command and control (C&C) addresses. This database is updated dynamically and stored on the FortiGate and requires a valid FortiGuard AntiVirus subscription.

When you block DNS requests to known botnet C&C addresses, using IPS, DNS lookups are checked against the botnet C&C database. All matching DNS lookups are blocked. Matching uses a reverse prefix match, so all sub-domains are also blocked.

To enable this feature, go to **Security Profiles > DNS Filter**, and enable **Block DNS requests to known botnet C&C**.

Static Domain Filter

The DNS **Static Domain Filter** allows you to block, exempt, or monitor DNS requests by using IPS to look inside DNS packets and match the domain being looked up with the domains on the static URL filter list. If there is a match the DNS request can be blocked, exempted, monitored, or allowed.

If blocked, the DNS request is blocked and so the user cannot look up the address and connect to the site.

If exempted, access to the site is allowed even if another method is used to block it.

CLI commands

- Rename `webfilter-sdns-server-ip` and `webfilter-sdns-server-port`:

```
config system fortiguard
    set sdns-server-ip x.x.x.x
    set sdns-server-port 53
end
```

- Configure DNS domain filter lists in order to decide access for specific domains:

```
config dnsfilter domain-filter
    edit {id}
        set id {integer}
        set name {string}
        set comment {string}
        config entries
            edit {id}
                set id {integer}
                set domain {string}
            next
        next
    next
end
```

```

        set type {simple | regex | wildcard}
        set action {block | allow | monitor}
        set status {enable | disable}
    next
next
end

```

- Configure DNS filter profile:

```

config dnsfilter profile
  edit "dns_profile1"
    set comment ''
    config domain-filter
      set domain-filter-table <id>
      set external-blocklist [addr1] [addr2] [addr3]
    end
    config ftgd-dns
      config filters
        edit 1
          set category 49
          set action block
          set log enable
        next
        edit 2
          set category 71
          set action monitor
          set log enable
        next
      end
    end
    set log-all-url disable
    set block-action redirect
    set redirect-portal 0.0.0.0
    set block-botnet enable
  next
end

```

- Configure DNS profile in a firewall policy:

```

config firewall policy
  edit 1
    set srcintf "any"
    set dstintf "any"
    set srcaddr "all"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "FTP"
    set utm-status enable
    set dnsfilter-profile "dns_profile1"
    set profile-protocol-options "default"
    set nat enable
  next
end

```

- Configure DNS profile in profile group:

```

config firewall profile-group
  edit "pgrp1"

```

```
set dnsfilter-profile "dns_profile1"
set profile-protocol-options "default"
next
end
```

DNS profile supports safe search

Users can take advantage of pre-defined DNS filter rules to edit DNS profiles and provide safe search for Google, Bing, and YouTube.

To add safe search to a DNS profile - GUI

1. Go to **Security Profiles > DNS Filter**.
2. Edit the default filter or create a new one.
3. Enable **Enforce 'Safe Search' on Google, Bing, YouTube**.
4. Select **Strict** or **Moderate** level for **Restrict YouTube Access**.

To add safe search to a DNS profile - CLI

```
config dnsfilter profile
edit "default"
set safe-search enable
set youtube-restrict {strict | moderate} (only available if safe-search enabled)
next
end
```

FortiGuard botnet protection

Preventing botnets from controlling your system is achieved by detecting and blocking connection attempts to known botnets. This feature also blocks connections to known phishing sites. The FortiGuard database is continually updated with addresses of known Command and Control (C&C) sites that botnet clients attempt to connect to, as well as addresses of known phishing URLs.

To enable botnet and phishing protection in a DNS Filter profile, enable **Block DNS requests to known botnet C&C**.

The latest botnet database is available from FortiGuard. To see the version of the database and display its contents, go to **System > FortiGuard > AntiVirus** and view the lists for **Botnet IPs** and **Botnet Domains**. You can look up more details about Botnet IPs and Domains on the [FortiGuard site](#).

You can block, monitor, or allow outgoing connections to botnet sites for each FortiGate interface.



The DNS Filter security profile and the botnet protection features are available for both proxy-based and flow-based inspection modes.

Application control

Using the Application Control Security Profiles feature, your FortiGate unit can detect and take action against network traffic depending on the application generating the traffic. Based on FortiGate Intrusion Protection protocol decoders, application control is a user-friendly and powerful way to use Intrusion Protection features to log and manage the behavior of application traffic passing through the FortiGate unit. Application control uses IPS protocol decoders that can analyze network traffic to detect application traffic even if the traffic uses non-standard ports or protocols. Application control supports detection for traffic using HTTP protocol (versions 1.0, 1.1, and 2.0).

The FortiGate unit can recognize the network traffic generated by a large number of applications. You can create application control sensors that specify the action to take with the traffic of the applications you need to manage and the network on which they are active, and then add application control sensors to the firewall policies that control the network traffic you need to monitor.

Fortinet is constantly adding to the list of applications detected through maintenance of the FortiGuard Application Control Database. This database is part of the FortiGuard Intrusion Protection System Database because intrusion protection protocol decoders are used for application control and both of these databases have the same version number.

Cloud Access Security Inspection (CASI) is merged with Application Control resulting in [changes to the GUI and the CLI](#).

You can identify the version of the application control database installed on your unit by going to the **Licenses** widget on the **Dashboard** and hovering over the **IPS & Application Control** line; the status, expiry date, and version will be displayed. Additionally, you can see the complete list of applications supported by FortiGuard Application Control on the [FortiGuard](#) site or <http://fortiguard.com/appcontrol>. This web page lists all of the supported applications. You can select any application name to see details about the application.



Application Control is a standard part of any FortiCare support contract and the database for Application Control signatures is separate from the IPS database. However, botnet application signatures are still part of the IPS signature database since these are more closely related with security issues and less about application detection.

This Handbook chapter includes [Inside FortiOS: Application Control](#) and provides readers an overview of the features and benefits of key FortiOS 5.6 components. For readers needing to delve into greater detail, we provide the following topics:

Application control concepts

You can control network traffic generally by the source or destination address, or by the port, the quantity or similar attributes of the traffic itself in the security policy. If you want to control the flow of traffic from a specific application, these methods may not be sufficient to precisely define the traffic. To address this problem, the application control feature examines the traffic itself for signatures unique to the application generating it. Application control does not require knowledge of any server addresses or ports. The FortiGate unit includes signatures for over 2,000 applications, services, and protocols.

Updated and new application signatures are delivered to your FortiGate unit as part of your FortiGuard Application Control Service subscription, which is a free service. Fortinet is constantly increasing the number of applications that this feature can detect by adding applications to the FortiGuard Application Control Database. Because intrusion protection protocol decoders are used for application control, the application control database is part of the FortiGuard Intrusion Protection System Database. Both of these databases have the same version number.

You can find the version of the application control database installed on your unit by going to the **Licenses** widget on the **Dashboard** and hovering over the **IPS& Application Control** line; the status, expiry date, and version will be displayed.

To see the complete list of applications supported by FortiGuard Application Control go to the [FortiGuard](http://fortiguard.com/appcontrol) site or <http://fortiguard.com/appcontrol>. This web page lists all of the supported applications. You can select any application name to see details about the application.

Enabling application control in profile-based modes

Application control examines your network traffic for traffic generated by the applications you want it to control. The configuration steps outlined below are for FortiGate's operating in proxy-based inspection and flow-based inspection with profile-based NGFW modes. For FortiGate's operating in NGFW policy-based mode, see [Enabling application control in NGFW policy-based mode](#).

General configuration steps

Follow the configuration procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. Create an application sensor.
2. Configure the sensor to include the signatures for the application traffic you want the FortiGate unit to detect.
3. Enable any other applicable options.
4. Enable application control in a security policy and select the application sensor.

Creating an application sensor

You need to create an application sensor before you can enable application control.

To create an application sensor

1. Go to **Security Profiles > Application Control**.
2. Select the **Create New** icon in the title bar of the **Edit Application Sensor** window.
3. In the **Name** field, enter the name of the new application sensor.
4. Optionally, enter descriptive **Comments**.

Adding applications to an application sensor

Once you have created an application sensor, you need to define the applications that you want to control. You can add applications and filters using categories, application overrides, and/or filter overrides. Categories will allow you to choose groups of signatures based on a category type. Application overrides allow you to choose individual applications. Filter overrides allow you to select groups of applications and override the application signature settings for them.

To add a category of signatures to the sensor.

1. Go to **Security Profiles > Application Control**.
2. Under **Categories**, you may select from the following:
 - Business
 - Cloud,.IT
 - Collaboration
 - Email
 - Game
 - General.Interest
 - Industrial
 - Mobile
 - Network.Service
 - P2P
 - Proxy
 - Remote.Access
 - Social.Media
 - Storage.Backup
 - Update
 - Video/Audio
 - VoIP
 - Web.Client
 - Unknown Applications

When selecting the category that you intend to work with, left click on the icon next to the category name to see a drop down menu that includes these actions:

- Allow
- Monitor
- Block
- Quarantine
- View Signatures

These actions are briefly defined under [Application control actions on page 2489](#).

3. If you wish to add individual applications, select **Add Signatures** under **Application Overrides**.
 - a. Use the **Add Filter** search field to narrow down the list of possible signatures by a series of attributes.
 - b. When finished, select **Use Selected Signatures**.
4. If you wish to add advanced filters, select **Add Filter** under **Filter Overrides**.
 - a. Use the Add Filter search field to narrow down the list of possible filters by a series of attributes.
 - b. When finished, select **Use Filters**.
4. Select, if applicable, from the following options:
 - **Allow and Log DNS Traffic**
 - **Replacement Messages for HTTP-based Applications**
6. Select **OK**.

Applying the application sensor to a security policy

An application sensor directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When an application sensor is selected in a security policy, its settings are applied to all the traffic the security policy handles.

To select the application sensor in a security policy — GUI

1. Go to **Policy & Objects > IPv4 Policy**.
2. Select a policy.
3. Click the **Edit** icon.
4. Under the heading **Security Profiles** toggle the button next to **Application Control** to enable the feature.
5. In the drop down menu field next to the **Application Control** select the application sensor you wish to apply to the policy.
6. Select **OK**.

Creating a new custom application signature

If you have to deal with an application that is not already in the **Application List** you have the option to create a new application signature.

1. Go to **Security Profiles > Application Control**.
2. Select the link in the upper right corner, **[View Application Signatures]**
3. Select the **Create New** icon
4. Give the new signature a name (no spaces) in the **Name** field.
5. Enter a brief description in the **Comments** field
6. Enter the text for the signature in the signature field. Use the rules found under [Custom IPS signature](#) to determine syntax.
7. Select **OK**.



You can configure rate based application control signatures in the CLI Console using similar IPS signature rate CLI commands.

For more information on this and the CLI syntax, see [IPS signature rate count threshold](#) on page 2505

Messages in response to blocked applications

Once an Application Control sensor has been configured to block a specified application and applied to a policy it would seem inevitable that at some point an application will end up getting blocked, even if it is only to test the functionality of the control. When this happens, the sensor can be set to either display a message to offending user or to just block without any notification. The default setting is to display a message. Setting this up is done in the CLI.

```
config application list
  edit <name of the sensor>
    set app-replacemsg {enable | disable}
  end
```



When blocking applications, there is no replacement message for SSL traffic with certificate inspection applied.

When SSL deep inspection is enabled, a replacement message will appear depending on the protocol. For example, with HTTP2, the blocking is done in the SSL key exchange once the first server packet is delivered and replacement messages can not be displayed.

P2P application detection

P2P software tends to be evasive. You may be able to enhance P2P application detection by matching patterns found in the most recent three minutes of P2P traffic to determine if new traffic is P2P. Three minutes is the length of time information about matched P2P traffic remains in shared memory.

For example, the CLI commands below will result in the Intrusion Prevention System (IPS) looking for patterns formed by Skype traffic.

```
config application list
  edit <app_list_str>
    set p2p-black-list skype
  end
end
```

Application control actions

Allow

This action allows the targeted traffic to continue on through the FortiGate unit.

Monitor

This action allows the targeted traffic to continue on through the FortiGate unit but logs the traffic for analysis.

Block

This action prevents all traffic from reaching the application and logs all occurrences.

Quarantine

This action allows you to quarantine or block access to an application for a specified duration that can be entered in days, hours, and minutes. The default is 5 minutes.

View Signatures

This option brings up a window that displays a list of the signatures with the following columns:

- Name
- Category
- Technology - Technology is broken down into 3 technology models as well as the more basic Network-Protocol which would can be used as a catch all for anything not covered by the more narrowly defined technologies of:
 - Browser-Based
 - Client-Server

- Peer-to-Peer
- Popularity - Popularity is broken down into 5 levels of popularity represented by stars.
- Risk - The Risk property does not indicate the level of risk but the type of impact that is likely to occur by allowing the traffic from that application to occur.

Traffic Shaping

Prior to the release of FortiOS 5.4.0, application control traffic shaping was configured in the **Security Profiles > Application Control** interface. There is now a specific section for traffic shaping policies in **Policy & Objects > Traffic Shaping Policy**. See [Traffic shaping methods](#) in the chapter on Traffic Shaping for details

Application considerations

Some applications behave differently from most others. You should be aware of these differences before using application control to regulate their use.

IM applications

IM applications are controlled by either permitting or denying the users from logging in to the service. Individual IM accounts are configured as to whether or not they are permitted and then there is a global policy for how to action unknown users, by the application and whether to add the user to the black list or the white list. IM applications fall under the **Collaboration** category in the application signature database.

Skype

Based on the NAT firewall type, Skype takes advantage of several NAT firewall traversal methods, such as STUN (Simple Traversal of UDP through NAT), ICE (Interactive Connectivity Establishment) and TURN (Traversal Using Relay NAT), to make the connection.

The Skype client may try to log in with either UDP or TCP, on different ports, especially well-known service ports, such as HTTP (80) and HTTPS (443), because these ports are normally allowed in firewall settings. A client who has previously logged in successfully could start with the known good approach, then fall back on another approach if the known one fails.

The Skype client could also employ Connection Relay. This means if a reachable host is already connected to the Skype network, other clients can connect through this host. This makes any connected host not only a client but also a relay server.

SPDY

SPDY (pronounced speedy, it's a trademarked name not an acronym) is a networking protocol developed to increase the speed and security of HTML traffic. It was developed primarily by Google. The Application Control engine recognizes this protocol and its required SSL/TLS component within Application Control sensors. It is counted as part of application traffic for Google and other sources that use the protocol.

Application control monitor

The application monitor enables you to gain insight into the applications generating traffic on your network. When monitor is enabled in an application sensor entry and that security profile is selected in a security policy, all the detected traffic required to populate the selected charts is logged to the SQL database on the FortiGate unit hard drive. The charts are available for display in the **Applications** section of the **FortiView** menu.



Because the application monitor relies on an SQL database, the feature is available only on FortiGate units with an internal hard drive.

Application monitor data is stored on the hard drive and restarting the system does not affect the stored monitor data.

Application control data is available in **Log & Report**, if enabled.

Application control examples

The scenarios below provide a better understanding of how to implement Application Control and give some ideas as to why it would be used.

- [Blocking instant messaging](#)
- [Allowing only software updates](#)
- [Blocking Windows XP with a custom signature](#)

Blocking instant messaging

Instant messaging use is not permitted at the Example Corporation. Application control helps enforce this policy.

The configuration steps outlined below are for FortiGate's operating in proxy-based inspection and flow-based inspection with profile-based NGFW modes. For FortiGate's operating in NGFW policy-based mode, see [Enabling application control in NGFW policy-based mode](#).

Steps in this process

1. First you will create an application sensor with a single entry that monitors the category that includes instant messaging applications. You will set the list action to **Monitor**.
2. Next you will assign the sensor to a policy.
3. Then you will identify the IM applications being used on your network and modify the application sensor to **Block** use of those messaging applications

To create the application sensor

1. Go to **Security Profiles > Application Control**.
2. Select the **Create New** icon in the title bar of the **Edit Application Sensor** window.
3. In the **Name** field, enter `no_IM` for the application sensor name.
4. If the **Collaboration** category is not already set to **Monitor**, then left-click on the icon next to that category and select **Monitor** from the dropdown menu.
5. Select **OK** to save the new sensor.

To enable application control and select the application sensor

1. Go to **Policy & Objects > IPv4 Policy**.
2. Select the security policy that allows the network users to access the Internet and choose **Edit**.
3. Under the heading **Security Profiles** toggle the button next to **Application Control** to turn it on.

4. In the drop down menu field next to the **Application Control** select the **no_IM** application sensor.
5. To inspect all traffic, **SSL/SSH inspection** must be set to **deep-inspection** profile.
6. Select **OK**.

To identify IM applications in use on your network

1. Go to **FortiView > Applications**.
2. Select a time period from the options in the upper-right corner of the window and examine the list of applications.
3. Identify any IM applications you wish to block.

To block IM applications in use on your network

1. Go to **Security Profiles > Application Control** and edit the **no_IM** application sensor.
2. Under **Application Overrides**, click on **Add Signatures**.
3. Filter by **Name** and select the IM applications you wish to block.
4. Click on **Use Selected Signatures**.

The selected application will appear under **Application Overrides** and the action will be set to **Block**.

5. Select **Apply**.

The IM applications identified will be blocked by the security policy that has the **no IM** application sensor applied to it. If other firewall policies handle traffic that users could use for applications in the same category, enable application control with the **no IM** application sensor for those policies as well.

Allowing only software updates

Some departments at Example Corporation do not require access to the Internet to perform their duties. Management therefore decided to block their Internet access. Software updates quickly became an issue because automatic updates will not function without Internet access and manual application of updates is time-consuming.

The solution is configuring application control to allow only automatic software updates to access the Internet.

The configuration steps outlined below are for FortiGate's operating in proxy-based inspection and flow-based inspection with profile-based NGFW modes. For FortiGate's operating in NGFW policy-based mode, see [Enabling application control in NGFW policy-based mode](#).

To create an application sensor — GUI

1. Go to **Security Profiles > Application Control**.
2. Select the **Create New** icon in the title bar of the **Edit Application Sensor** window.
3. In the **Name** field, enter `Updates_Only` as the application sensor name.
4. Using the left-click and drop down on the items in the **Category** list.
 - a. Select **Monitor** from the dropdown menu.
 - b. Select **Block** for the rest of the categories.
5. Select **OK**.

To create an application sensor — CLI

```
config application list
```

```

edit Updates_Only
  config entries
    edit 1
      set category 17
      set action pass
    end
  set other-application-action block
  set unknown-application-action block
end

```



You will notice that there are some differences in the naming convention between the GUI and the CLI. For instance the **Action** in the CLI is “pass” and the **Action** in the GUI is “**Monitor**”.

Selecting the application sensor in a security policy

An application sensor directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When an application sensor is selected in a security policy, its settings are applied to all the traffic the security policy handles.

To select the application sensor in a security policy — GUI

1. Go to **Policy & Objects > IPv4 Policy**.
2. Select a policy.
3. Select the **Edit** icon.
4. Under the heading **Security Profiles** toggle the button next to **Application Control** to turn it on.
5. In the drop down menu field next to the **Application Control** select the `Updates_only` list.
6. Select **OK**.

To select the application sensor in a security policy — CLI

```

config firewall policy
  edit 1
    set utm-status enable
    set profile-protocol-options default
    set application-list Updates_Only
  end

```

Traffic handled by the security policy you modified will be scanned for application traffic. Software updates are permitted and all other application traffic is blocked.

Blocking Windows XP with a custom signature

In this example, you will use application control to block web traffic from PCs running Windows operating systems NT 5, including Windows XP and Windows Server 2003 (includes Windows virtual machines).

When a computer's operating system lacks vendor support, it becomes a threat to the network because newly discovered exploits will not be patched. Using the FortiGate application control feature, you can restrict these computers from accessing external resources.

This example will only block web traffic from computers running the affected operating systems. If you wish to block these computers from being on the network entirely, further action will be necessary. However, the logs generated can be used to identify the computers you wish to block.

1. Go to **System > Feature Select**. Enable **Application Control** and **Apply** your changes.
2. Go to **Security Profiles > Application Control** and select **View Application Signatures**.
3. Create a new signature with the syntax below. You can copy and paste the text into the **Signature** field. Name the signature *Block-Windows-NT5*.

```
F-SBID(--attack_id 8055;--vuln_id 8055;--name
"Windows.NT.5.Web.Surfing";--flow from_client;--pattern !"FCT";--
pattern "Windows NT 5.";--no_case;--context header;--weight 40;--
service HTTP;--protocol tcp;--app_cat 25;--default_action drop_
session;)
```

If you do not include keyword / value pairs for `--attack_id` or `--vuln_ID` in the signature, the FortiGate will automatically assign values.

The signature will appear at the top of the application list and be listed in the **Web.Client** category.

4. Go to **Security Profiles > Application Control** and edit the **default** policy.
5. Under **Application Overrides**, select **Add Signatures**. The new signature should appear at the top of the list. If it does not, search for the signature's name.
6. Select the signature, then select **Use Selected Signatures**.
7. Go to **Policy & Objects > IPv4 Policy** and edit the policy that allows connections from the internal network to the Internet.
8. Under **Security Profiles**, turn on **Application Control** and use the **default** profile.

Results

When a PC running one of the affected operating systems attempts to connect to the Internet using a browser, a blocked message appears. Because Application Control uses flow-based inspection, if you apply an additional security profile to your traffic that is proxy-based, the connection will simply timeout rather than display the replacement message. However, Application Control will still function.

PCs running other operating systems, including later versions of Windows, are not affected.

Go to **FortiView > All Sessions** and select the **5 minutes** view.

Filter the results to show sessions that were blocked.

You will see that the Application Control signature, shown in the **Application Name** column, was used to block traffic from PCs running older Windows versions.

For further reading, see [Custom Application & IPS Signatures](#).

Intrusion prevention

The FortiOS Intrusion Prevention System (IPS) combines signature detection and prevention with low latency and excellent reliability. With intrusion protection, you can create multiple IPS sensors, each containing a complete configuration based on signatures. Then, you can apply any IPS sensor to any security policy.

This section describes how to configure the FortiOS Intrusion Prevention settings.

This Handbook chapter includes [Inside FortiOS: Intrusion Prevention System](#) providing readers an overview of the features and benefits of key FortiOS 5.6 components. For readers needing to delve into greater detail, we provide the following:

IPS concepts

The FortiOS Intrusion Prevention System (IPS) protects your network from outside attacks. Your FortiGate unit has two techniques to deal with these attacks: anomaly- and signature-based defense.

Anomaly-based defense

Anomaly-based defense is used when network traffic itself is used as a weapon. A host can be flooded with far more traffic than it can handle, making the host inaccessible. The most common example is the denial of service (DoS) attack, in which an attacker directs a large number of computers to attempt normal access of the target system. If enough access attempts are made, the target is overwhelmed and unable to service genuine users. The attacker does not gain access to the target system, but it is not accessible to anyone else.

The FortiGate DoS feature will block traffic above a certain threshold from the attacker and allow connections from other legitimate users. The DoS policy configuration can be found in the Firewall chapter of the Handbook.

Access control lists in DoS Policies

This feature allows you to define a list of IPs/subnets/ranges in a DoS policy, and block those IPs from sending any traffic, by way of an ACL (access control list). The ACL looks similar to a firewall policy, but only checks source IP, destination IP, destination port, and protocol. To configure in the GUI, go to **Policy & Objects > IPv4 Access Control List** and create a new policy. Enter the incoming interface, the source address, the destination address, the services impacted, and, optionally, enter a comment.

CLI Syntax

```
config firewall acl
  edit 1
    set interface "port1"
    set srcaddr "google-drive"
    set dstaddr "all"
    set service "ALL"
  next
end
```

Signature-based defense

Signature-based defense is used against known attacks or vulnerability exploits. These often involve an attacker attempting to gain access to your network. The attacker must communicate with the host in an attempt to gain access and this communication will include particular commands or sequences of commands and variables. The IPS signatures include these command sequences, allowing the FortiGate unit to detect and stop the attack.

Signatures

IPS signatures are the basis of signature-based intrusion prevention. Every attack can be reduced to a particular string of commands or a sequence of commands and variables. Signatures include this information so your FortiGate unit knows what to look for in network traffic.

Signatures also include characteristics about the attack they describe. These characteristics include the network protocol in which the attack will appear, the vulnerable operating system, and the vulnerable application.

To view the complete list of signatures, go to **Security Profiles > Intrusion Prevention**, and select **View IPS Signatures**. This will include the predefined signatures and any custom signatures that you may have created.

With the release of FortiOS 5.6, the IPS signatures list page shows which IPS package is currently deployed. Users can also change their IPS package by hovering over the information icon next to the IPS package name. Text will appear that links directly to the FortiGate's **System > FortiGuard** page from the IPS Signatures list page.

Protocol decoders

Before examining network traffic for attacks, the IPS engine uses protocol decoders to identify each protocol appearing in the traffic. Attacks are protocol-specific, so your FortiGate unit conserves resources by looking for attacks only in the protocols used to transmit them. For example, the FortiGate unit will only examine HTTP traffic for the presence of a signature describing an HTTP attack.

IPS engine

Once the protocol decoders separate the network traffic by protocol, the IPS engine examines the network traffic for the attack signatures.

IPS sensors

The IPS engine does not examine network traffic for all signatures. You must first create an IPS sensor and specify which signatures are included. Add signatures to sensors individually using signature entries, or in groups using IPS filters.

To view the IPS sensors, go to **Security Profiles > Intrusion Prevention**.

You can group signatures into IPS sensors for easy selection when applying to firewall policies. You can define signatures for specific types of traffic in separate IPS sensors, and then select those sensors in profiles designed to handle that type of traffic. For example, you can specify all of the web-server related signatures in an IPS sensor, and that sensor can then be applied to a firewall policy that controls all of the traffic to and from a web server protected by the unit.

The FortiGuard Service periodically updates the pre-defined signatures, with signatures added to counter new threats. Since the signatures included in filters are defined by specifying signature attributes, new signatures matching existing filter specifications will automatically be included in those filters. For example, if you have a

filter that includes all signatures for the Windows operating system, your filter will automatically incorporate new Windows signatures as they are added.

Each IPS sensor consists of two parts: filters and overrides. Overrides are always checked before filters.

Each filter consists of a number of signature attributes. All of the signatures with those attributes, and only those attributes, are checked against traffic when the filter is run. If multiple filters are defined in an IPS Sensor, they are checked against the traffic one at a time, from top to bottom. If a match is found, the unit takes the appropriate action and stops further checking.

A signature override can modify the behavior of a signature specified in a filter. A signature override can also add a signature not specified in the sensor's filters. Custom signatures are included in an IPS sensor using overrides.

The signatures in the overrides are first compared to network traffic. If the IPS sensor does not find any matches, it then compares the signatures in each filter to network traffic, one filter at a time, from top to bottom. If no signature matches are found, the IPS sensor allows the network traffic.

The signatures included in the filter are only those matching every attribute specified. When created, a new filter has every attribute set to **all** which causes every signature to be included in the filter. If the severity is changed to high, and the target is changed to server, the filter includes only signatures checking for high priority attacks targeted at servers.

IPS filters

IPS sensors contain one or more IPS filters. A filter is a collection of signature attributes that you specify. The signatures that have all of the attributes specified in a filter are included in the IPS filter.

For example, if your FortiGate unit protects a Linux server running the Apache web server software, you could create a new filter to protect it. By setting **OS** to **Linux**, and **Application** to **Apache**, the filter will include only the signatures that apply to both Linux and Apache. If you wanted to scan for all the Linux signatures and all the Apache signatures, you would create two filters, one for each.

To view the filters in an IPS sensor, go to **Security Profiles > Intrusion Prevention**, select the IPS sensor containing the filters you want to view, and select **Edit**.

Custom/predefined signature entries

Signature entries allow you to add an individual custom or predefined IPS signature. If you need only one signature, adding a signature entry to an IPS sensor is the easiest way. Signature entries are also the only way to include custom signatures in an IPS sensor.

Another use for signature entries is to change the settings of individual signatures that are already included in a filter within the same IPS sensor. Add a signature entry with the required settings above the filter, and the signature entry will take priority.

Policies

To use an IPS sensor, you must select it in a security policy or an interface policy. An IPS sensor that is not selected in a policy will have no effect on network traffic.

IPS is most often configured as part of a security policy. Unless stated otherwise, discussion of IPS sensor use will be in regards to firewall policies in this document.

Session timers for IPS sessions

A session time-to-live (TTL) timer for IPS sessions is available to reduce synchronization problems between the FortiOS Kernel and IPS, and to reduce IPS memory usage. The timeout values can be customized.

Enabling IPS scanning

Enabling IPS scanning involves two separate features of FortiOS 5.6:

- The security policy allows certain network traffic based on the sender, receiver, interface, traffic type, and time of day. Firewall policies can also be used to deny traffic, but those policies do not apply to IPS scanning.
- The IPS sensor contains filters, signature entries, or both. These specify which signatures are included in the IPS sensor.

When IPS is enabled and an IPS sensor is selected in a security policy, and all network traffic matching the policy will be checked for the signatures in the IPS sensor.

General configuration steps

For best results in configuring IPS scanning, follow the procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. Create an IPS sensor.
2. Add signatures and /or filters.
These can be:
 - Pattern based
 - Rate based
 - Customized
3. Select a security policy or create a new one.
4. In the security policy, turn on **IPS**, and choose the IPS sensor from the list.

All the network traffic controlled by this security policy will be processed according to the settings in the policy. These settings include the IPS sensor you specify in the policy.

Creating an IPS sensor

You need to create an IPS sensor before specific signatures or filters can be chosen. The signatures can be added to a new sensor before it is saved. However, it is good practice to keep in mind that the sensor and its included filters are separate things, and that they are created separately.

To create a new IPS sensor

1. Go to **Security Profiles > Intrusion Prevention**.
2. Select the **Create New** icon in the top of the Edit IPS Sensor window.
3. Enter the name of the new IPS sensor.
4. Optionally, enter a comment. The comment will appear in the IPS sensor list.
5. Select **OK**.

A newly created sensor is empty and contains no filters or signatures. You need to add one or more filters or signatures before the sensor will be of any use.

Adding an IPS filter to a sensor

While individual signatures can be added to a sensor, a filter allows you to add multiple signatures to a sensor by specifying the characteristics of the signatures to be added.

To create a new pattern based signature and filter

1. Go to **Security Profiles > Intrusion Prevention**.
2. Select the IPS sensor to which you want to add the filter using the drop-down list in the top row of the Edit IPS Sensor window or by going to the list window.
3. Under **IPS Filters**, select **Add Filter**.
4. Configure the filter that you require. Signatures matching all of the characteristics you specify in the filter will be included in the filter. Once finished, select **Use Filters**.

Application refers to the application affected by the attack and filter options include over 25 applications.

OS refers to the Operating System affected by the attack. The options include **BSD**, **Linux**, **MacOS**, **Other**, **Solaris**, and **Windows**.

Protocol refers to the protocol that is the vector for the attack; filter options include over 35 protocols, including "other."

Severity refers to the level of threat posed by the attack. The options include **Critical**, **High**, **Medium**, **Low**, and **Info**.

Target refers to the type of device targeted by the attack. The options include **client** and **server**.

5. Once you have selected the filters you wish to add, right-click the filters and choose an action for when a signature is triggered:

Action	Description
Pass	Select Pass to allow traffic to continue to its destination. Note: to see what the default for a signature is, go to the IPS Signatures page and enable the column Action , then find the row with the signature name in it.
Monitor	Select Monitor to allow traffic to continue to its destination and log the activity. The log will appear under Log & Report but will only be visible in the GUI in the event of an intrusion.
Block	Select Block to drop traffic matching any the signatures included in the filter.
Reset	Select Reset to reset the session whenever the signature is triggered. In the CLI this action is referred to as Reject.
Default	Select Default to use the default action of the signature.
Quarantine	The quarantine based on the attacker's IP Address - Traffic from the Attacker's IP address is refused until the expiration time from the trigger is reached. You may set the Quarantine Duration to any number of Days , Hours , or Minutes .

Action	Description
Packet Logging	<p>Select to enable packet logging for the filter.</p> <p>When you enable packet logging on a filter, the unit saves a copy of the packets that match any signatures included in the filter. The packets can be analyzed later.</p> <p>For more information about packet filtering, see "Configuring packet logging options".</p>

6. Select **Apply**.

The filter is created and added to the filter list.

Adding rate based signatures

These are a subset of the signatures that are found in the database that are normally set to monitor. This group of signatures is for vulnerabilities that are normally only considered a serious threat when the targeted connections come in multiples, a little like DoS attacks.

Adding a rate based signature is straight forward. Select the enable button in the Rate Based Signature table that corresponds with the desired signature.

Customized signatures

Customized signatures must be created before they can be added to the sensor. To get more details on customized signatures check the [Custom Application & IPS Signatures](#) chapter.

Updating predefined IPS signatures

The FortiGuard Service periodically updates the predefined signatures and adds new signatures to counter emerging threats as they appear.

To ensure that your system is providing the most protection available, these updates can be scheduled as often as on an hourly basis. To configure this feature, go to **System > FortiGuard**. Under **AntiVirus & IPS Updates**, enable **Scheduled Updates**. From here you can set the updates to occur on a consistent weekly, daily, or even hourly basis.

Because the signatures included in filters are defined by specifying signature attributes, new signatures matching existing filter specifications will automatically be included in those filters. For example, if you have a filter that includes all signatures for the Windows operating system, your filter will automatically incorporate new Windows signatures as they are added.

Viewing and searching predefined IPS signatures

Go to **Security Profiles > Intrusion Prevention**. Select **[View IPS Signatures]** to view the list of existing IPS signatures. You may find signatures by paging manually through the list, apply filters, or by using the search field.

Searching manually

Signatures are displayed in a paged list, with 50 signatures per page. The bottom of the screen shows the current page and the total number of pages. You can enter a page number and press enter, to skip directly to that page. Previous Page and Next Page buttons move you through the list, one page at a time. The First Page and Last Page button take you to the beginning or end of the list.

Searching CVE-IDs

A **CVE-ID** column displaying CVE-IDs can be optionally added to the IPS Signatures list, however the column is only available if the IPS package contains CVE-IDs for signatures. CVE-IDs can be numerically filtered by selecting the CVE-ID column's arrows.

Applying filters

You can enter criteria for one or more columns, and only the signatures matching all the conditions you specify will be listed.

To apply filters

1. Go to **Security Profiles > Intrusion Prevention**. Select **[View IPS Signatures]**.
2. Select column by which to filter.
3. Select the funnel/filter icon and enter the value or values to filter by.
4. Use additional columns as needed to refine search.

The available options vary by column. For example, **Enable** allows you to choose between two options, while OS has multiple options, and you may select multiple items together. Filtering by name allows you to enter a text string and all signature names containing the string will be displayed.

IPS processing in an HA cluster

IPS processing in an HA cluster is no different than with a single FortiGate unit, from the point of view of the network user. The difference appears when a secondary unit takes over from the primary, and what happens depends on the HA mode.

Active-passive

In an active-passive HA cluster, the primary unit processes all traffic just as it would in a stand-alone configuration. Should the primary unit fail, a secondary unit will assume the role of the primary unit and begin to process network traffic. By default, the state of active communication sessions are not shared with secondary units and will not survive the fail-over condition. Once the sessions are reestablished however, traffic processing will continue as normal.

If your network requires that active sessions are taken over by the new primary unit, select **Enable Session Pick-up** in your HA configuration. Because session information must be sent to all subordinate units on a regular basis, session pick-up is a resource-intensive feature and is not enabled by default.

Active-active

The fail-over process in an active-active cluster is similar to an active-passive cluster. When the primary unit fails, a secondary unit takes over and traffic processing continues. The load-balancing schedule used to distribute sessions to the cluster members is used by the new primary unit to redistribute sessions among the remaining subordinate units. If session pick-up is not enabled, the sessions active on the failed primary are lost, and the sessions redistributed among the secondary units may also be lost. If session pick-up is enabled, all sessions are handled according to their last-known state.

Configure IPS options

The following IPS configuration options are available:

- Malicious URL database for drive-by exploits detection
- Customizable replacement message when IPS blocks traffic
- Hardware acceleration
- Extended IPS database
- Configuring the IPS engine algorithm
- Configuring the IPS engine-count
- Configuring fail-open
- Configuring the session count accuracy
- Configuring IPS intelligence
- Configuring the IPS buffer size
- Configuring protocol decoders
- Configuring security processing modules
- IPS signature rate count threshold
- Geographic location filter

Malicious URL database for drive-by exploits detection

This feature uses a local malicious URL database on the FortiGate to assist in drive-by exploits detection. The database contains all malicious URLs active in the last one month, and all drive-by exploit URLs active in the last three months. The number of URLs controlled are in the one million range.

```
config ips sensor
  edit <profile>
    set block-malicious-url [enable | disable]
  next
end
```

Customizable replacement message when IPS blocks traffic

You can edit a replacement message that will appear specifically for IPS sensor blocked Internet access. Go to **System > Replacement Messages**, select **Extended View** and find **IPS Sensor Block Page** under the **Security** heading.

Hardware acceleration for flow-based security profiles (NTurbo and IPSA)

Some FortiGate models support a feature call NTurbo that can offload flow-based firewall sessions to NP4 or NP6 network processors. Some FortiGate models also support offloading enhanced pattern matching for flow-based security profiles to CP8 or CP9 content processors. You can use the following command to configure NTurbo and IPSA:

```
config ips global
  set np-accel-mode {none | basic}
  set cp-accel-mode {none | basic | advanced}
end
```

If the `np-accel-mode` option is available, your FortiGate supports NTurbo: `none` disables NTurbo and `basic` (the default) enables NTurbo. If the `cp-accel-mode` option is available your FortiGate supports IPSA: `none` disables IPSA, `basic` enables basic IPSA and `advanced` enables enhanced IPSA which can offload more types of pattern matching than basic IPSA. `advanced` is only available on FortiGate models with two or more CP8 processors or one or more CP9 processors.

See the **Hardware Acceleration** handbook chapter for more information about NTurbo and IPSA.

Extended IPS database

Some models have access to an extended IPS Database. The extended database may affect the performance of the FortiGate unit so depending on the model of the FortiGate unit the extended database package may not be enabled by default. For example, the D-series Desktop model have this option disabled by default.

This feature can only be enabled through the CLI.

```
config ips global
    set database extended
end
```

Configuring the IPS engine algorithm

The IPS engine is able to search for signature matches in two ways. One method is faster but uses more memory, the other uses less memory but is slower. Use the `algorithm` CLI command to select one method:

```
config ips global
    set algorithm {super | high | low | engine-pick}
end
```

Specify `high` to use the faster more memory intensive method or `low` for the slower memory efficient method. The setting `super` improves the performance for FortiGate units with more than 4GB of memory. The default setting is `engine-pick`, which allows the IPS engine to choose the best method on the fly.

Configuring the IPS engine-count

FortiGate units with multiple processors can run more than one IPS engine concurrently. The `engine-count` CLI command allows you to specify how many IPS engines are used at the same time:

```
config ips global
    set engine-count <int>
end
```

The recommended and default setting is 0, which allows the FortiGate unit to determine the optimum number of IPS engines.

Configuring fail-open

IPS is likely more important to your network than uninterrupted flow of network traffic, so the fail-open behaviour of the IPS engine is disabled by default. If you would like to enable the fail-open option, use the following syntax. When enabled, if the IPS engine fails for any reason, it will fail open. This applies for inspection of all the protocols inspected by FortiOS IPS protocol decoders, including but not limited to HTTP, HTTPS, FTP, SMTP, POP3, IMAP, etc. This means that traffic continues to flow without IPS scanning. To enable:

```
config ips global
    set fail-open {enable | disable}
end
```

The default setting is `disable`.

Configuring the session count accuracy

The IPS engine can keep track of the number of open session in two ways. An accurate count uses more resources than a less accurate heuristic count.

```
config ips global
    set session-limit-mode {accurate | heuristic}
```

```
end
```

The default is `heuristic`.

Configuring IPS intelligence

Starting with FortiOS 5.2, `intelligent-mode` is a new adaptive detection method. This command is enabled the default and it means that the IPS engine will perform adaptive scanning so that, for some traffic, the FortiGate can quickly finish scanning and offload the traffic to NPU or kernel. It is a balanced method which could cover all known exploits. When disabled, the IPS engine scans every single byte.

```
config ips global
    set intelligent-mode {enable|disable}
end
```

Configuring the IPS buffer size

Set the size of the IPS buffer.

```
config ips global
    set socket-size <int>
end
```

The acceptable range is from 1 to 64 megabytes. The default size varies by model. In short, `socket-size` determines how much data the kernel passes to the IPS engine each time the engine samples packets.

Configuring protocol decoders

The FortiGate Intrusion Prevention system uses protocol decoders to identify the abnormal traffic patterns that do not meet the protocol requirements and standards. For example, the HTTP decoder monitors traffic to identify any HTTP packets that do not meet the HTTP protocol standards.

To change the ports a decoder examines, you must use the CLI. In this example, the ports examined by the DNS decoder are changed from the default 53 to 100, 200, and 300.

```
config ips decoder dns_decoder
    set port_list "100,200,300"
end
```

You cannot assign specific ports to decoders that are set to **auto** by default. These decoders can detect their traffic on any port. Specifying individual ports is not necessary.

Configuring security processing modules

FortiGate Security Processing Modules, such as the CE4, XE2, and FE8, can increase overall system performance by accelerating some security and networking processing on the interfaces they provide. They also allow the FortiGate unit to offload the processing to the security module, thereby freeing up its own processor for other tasks. The security module performs its own IPS and firewall processing, but you can configure it to favor IPS in hostile high-traffic environments.

If you have a security processing module, use the following CLI commands to configure it to devote more resources to IPS than firewall. This example shows the CLI commands required to configure a security module in slot 1 for increased IPS performance.

```
config system amc-slot
    edit sw1
        set optimization-mode fw-ips
    end
end
```

```

set ips-weight balanced
set ips-p2p disable
set ips-fail-open enable
set fp-disable none
set ipsec-inb-optimization enable
set syn-proxy-client-timer 3
set syn-proxy-server-timer 3
end

```

In addition to offloading IPS processing, security processing modules provide a hardware accelerated SYN proxy to defend against SYN flood denial of service attacks. When using a security module, configure your DoS anomaly check for `tcp_syn_flood` with the **Proxy** action. The **Proxy** action activates the hardware accelerated SYN proxy.

IPS signature rate count threshold

The IPS signature threshold can allow configuring a signature so that it will not be triggered until a rate count threshold is met. This provides a more controlled recording of attack activity. For example, if multiple login attempts produce a failed result over a short period of time then an alert would be sent and perhaps traffic blocked. This would be a more rational response than sending an alert every time a login failed.

The syntax for this configuration is as follows:

```

config ips sensor
  edit default
    config entries
      edit <Filter ID number>
        set rule <*id>
        set rate-count <integer between 1 - 65535>
        set rate-duration <integer between 1 - 65535>

```

The value of the `rate-duration` is an integer for the time in seconds.

```
set rate-mode <continuous | periodical>
```

The `rate-mode` refers to how the count threshold is met.

If the setting is “continuous”, and the action is set to block, as soon as the `rate-count` is reached the action is engaged. For example, if the count is 10, as soon as the signature is triggered 10 times the traffic would be blocked.

If the setting is “periodical”, the FortiGate allows up to the value of the `rate-count` incidents where the signature is triggered during the `rate-duration`. For example, if the rate count is 100 and the duration is 60, the signature would need to be triggered 100 times in 60 seconds for the action to be engaged.

```
set rate-track <dest-ip | dhcp-client-mac | dns-domain | none | src-ip>
```

This setting allows the tracking of one of the protocol fields within the packet.

Geographic location filter

Place filters based on geographical location. Note that routes will not be installed if the resolved IPv6 address belongs to the country in the filter.

Any country entered for `geo-filter` will prevent all destination addresses that belong to that country from being installed into static routing table:

```
config webfilter {ips-urlfilter-setting | ips-urlfilter-setting6}
```



```
edit <address>
    set geo-filter <country-name>
next
end
```

Use the following diagnose command to list the IPv4 and/or IPv6 IP ranges of a specific country:

```
diagnose geoip {iprange6 | iprange} <country-name>
```

Enabling IPS packet logging

Packet logging saves the network packets containing the traffic matching an IPS signature to the attack log. The FortiGate unit will save the logged packets to wherever the logs are configured to be stored, whether memory, internal hard drive, a FortiAnalyzer unit, or the FortiGuard Analysis and Management Service.

You can enable packet logging in the filters. Use caution in enabling packet logging in a filter. Filters configured with few restrictions can contain thousands of signatures, potentially resulting in a flood of saved packets. This would take up a great deal of space, require time to sort through, and consume considerable system resources to process. Packet logging is designed as a focused diagnostic tool and is best used with a narrow scope.



Although logging to multiple FortiAnalyzer units is supported, packet logs are not sent to the secondary and tertiary FortiAnalyzer units. Only the primary unit receives packet logs.

To enable packet logging for a filter

1. Create a filter in an IPS sensor.
2. After creating the filter, right-click the filter, and select **Enable** in the **Packet Logging** column of the filter table.
3. Select the IPS sensor in the security policy that allows the network traffic the FortiGate unit will examine for the signature.

For information on viewing and saving logged packets, see [Configuring packet logging options](#) below.

IPS logging changes

IPS operations severely affected by disk logging are moved out of the quick scanning path, including logging, SNMP trap generation, quarantine, etc.

Scanning processes are dedicated to nothing but scanning, which results in more evenly distributed CPU usage. Slow (IPS) operations are taken care of in a dedicated process, which usually stays idle.



Setting `packet-log-history` to a value larger than 1 can affect the performance of the FortiGate unit because network traffic must be buffered. The performance penalty depends on the model, the setting, and the traffic load.

Other IPS examples

Configuring basic Intrusion Prevention

Small offices, whether they are small companies, home offices, or satellite offices, often have very simple needs. This example details how to enable IPS on a FortiGate unit located in a satellite office. The satellite office

contains only Windows clients.

Creating an IPS sensor

Most IPS settings are configured in an IPS sensor. IPS sensors are selected in firewall policies. This way, you can create multiple IPS sensors, and tailor them to the traffic controlled by the security policy in which they are selected. In this example, you will create one IPS sensor.

To create an IPS sensor— GUI

1. Go to **Security Profiles > Intrusion Prevention**.
2. Select the **Create New** icon in the top of the Edit IPS Sensor window.
3. In the **Name** field, enter `basic_ips`.
4. In the **Comments** field, enter `IPS for Windows clients`.
5. Select **OK**.
6. Select the **Create New** drop-down to add a new component to the sensor and for the **Sensor Type** choose **Filter Based**.
7. In the Filter Options choose the following:
 - a. For **Severity**: select all of the options
 - b. For **Target**: select **Client** only.
 - c. For **OS**: select **Windows** only.
8. For the **Action** leave as the default.
9. Select **OK** to save the filter.
10. Select **OK** to save the IPS sensor.

To create an IPS sensor — CLI

```
config ips sensor
  edit basic_ips
    set comment "IPS for Windows clients"
    config entries
      edit 1
        set location client
        set os windows
      end
    end
  end
end
```

Selecting the IPS sensor in a security policy

An IPS sensor directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When an IPS sensor is selected in a security policy, its settings are applied to all the traffic the security policy handles.

To select the IPS sensor in a security policy — GUI

1. Go to **Policy & Objects > IPv4 Policy**.
2. Select a policy.
3. Select the **Edit** icon.
4. Enable the **IPS** option under **Security Profiles**.

5. Select the preferred IPS sensor from the dropdown menu.
6. Select **OK** to save the security policy.

To select the IPS sensor in a security policy — CLI

```
config firewall policy
  edit 1
    set utm-status enable
    set ips-sensor basic_ips
  end
```


The IPS sensor in this example is `basic_ips`. All traffic handled by the security policy you modified will be scanned for attacks against Windows clients. A small office may have only one security policy configured. If you have multiple policies, consider enabling IPS scanning for all of them.

Using IPS to protect your web server

Many companies have web servers and they must be protected from attack. Since web servers must be accessible, protection is not as simple as blocking access. IPS is one tool your FortiGate unit has to allow you to protect your network.

In this example, we will configure IPS to protect a web server. As shown below, a FortiGate unit protects a web server and an internal network. The internal network will have its own policies and configuration but we will concentrate on the web server in this example.

A simple network configuration



The FortiGate unit is configured with:

- a virtual IP to give the web server a unique address accessible from the Internet.
- a security policy to allow access to the web server from the Internet using the virtual IP.

To protect the web server using intrusion prevention, you need to create an IPS sensor, populate it with filters, then enable IPS scanning in the security policy.

To create an IPS sensor

1. Go to **Security Profiles > Intrusion Prevention**.
2. Select **Create New**.
3. Enter `web_server` as the name of the new IPS sensor.
4. Select **OK**.

The new IPS sensor is created but it has no filters, and therefore no signatures are included.

The web server operating system is Linux, so you need to create a filter for all Linux server signatures.

To create the Linux server filter

1. Go to **Security Profiles > Intrusion Prevention**.
2. Select the `web_server` IPS sensor and select the **Edit** icon.
3. In the **Pattern Based Signatures and Filters** section, select **Create New**.
4. For **Sensor Type**, select **Filter Based**.
5. For **Filter Options**.
6. In the Filter Options choose the following:
 - a. For **Severity**: select all of the options
 - b. For **Target**: select **server** only.
 - c. For **OS**: select **Linux** only.
7. Select **OK**.

The filter is saved and the IPS sensor page reappears. In the filter list, find the **Linux Server** filter and look at the value in the **Count** column. This shows how many signatures match the current filter settings. You can select the **View Rules** icon to see a listing of the included signatures.

To edit the security policy

1. Go to **Policy & Objects > IPv4 Policy** select security policy that allows access to the web server, and select the **Edit** icon.
2. Enable IPS option and choose the `web_server` IPS sensor from the list.
3. Select **OK**.

Since IPS is enabled and the `web_server` IPS sensor is specified in the security policy controlling the web server traffic, the IPS sensor examines the web server traffic for matches to the signatures it contains.

Create and test a packet logging IPS sensor

In this example, you create a new IPS sensor and include a filter that detects the EICAR test file and saves a packet log when it is found. This is an ideal first experience with packet logging because the EICAR test file can cause no harm, and it is freely available for testing purposes.

Create an IPS sensor

1. Go to **Security Profiles > Intrusion Prevention**.
2. Select **Create New**.
3. Name the new IPS sensor `EICAR_test`.
4. Select **OK**.

Create an entry

1. Select the **Create New**.
2. For **Sensor Type** choose **Specify Signatures**.
3. Rather than search through the signature list, use the name filter by selecting the search icon over the header of the **Signature** column.
4. Enter `EICAR` in the Search field.
5. Highlight the `Eicar.Virus.Test.File` signature by clicking on it.
6. Select **Block** as the **Action** for the `EICAR test` sensor in the **IPS Signatures** table.
7. Enable **Packet Logging**.
8. Select **OK** to save the IPS sensor.

Add the IPS sensor to the security policy allowing Internet access

1. Go to **Policy & Objects > IPv4 Policy**.
2. Select the security policy that allows you to access the Internet.
3. Select the **Edit** icon.
4. Go to **Security Profiles** and enable **IPS** and choose `EICAR test` from the available IPS sensors..
5. Enable **Log Allowed Traffic** and select **All Sessions**.
6. Select **OK**.

With the IPS sensor configured and selected in the security policy, the FortiGate unit blocks any attempt to download the EICAR test file.

Test the IPS sensor

1. Using your web browser, go to http://www.eicar.org/anti_virus_test_file.htm.
2. Scroll to the bottom of the page and select **eicar.com** from the row labeled as using the standard HTTP protocol.
3. The browser attempts to download the requested file and,
 - If the file is successfully downloaded, the custom signature configuration failed at some point. Check the custom signature, the IPS sensor, and the firewall profile.
 - If the download is blocked with a high security alert message explaining that you're not permitted to download the file, the EICAR test file was blocked by the FortiGate unit antivirus scanner before the IPS sensor could examine it. Disable antivirus scanning and try to download the EICAR test file again.
 - If no file is downloaded and the browser eventually times out, the custom signature successfully detected the EICAR test file and blocked the download.

Viewing the packet log

1. Go to **Log & Report > Forward Traffic**.
2. Locate the log entry that recorded the blocking of the EICAR test file block. The Message field data will be `tools: EICAR.AV.Test.File.Download`.
3. Select the **View Packet Log** icon in the **Packet Log** column.
4. The packet log viewer is displayed.

Configuring a Fortinet Security Processing module

The Example Corporation has a web site that is the target of SYN floods. While they investigate the source of the attacks, it's very important that the web site remain accessible. To enhance the ability of the company's FortiGate-100D to deal with SYN floods, the administrator will install an ASM-CE4 Fortinet Security Processing module and have all external access to the web server come through it.

The security processing modules not only accelerate and offload network traffic from the FortiGate unit's processor, but they also accelerate and offload security and content scanning. The ability of the security module to accelerate IPS scanning and DoS protection greatly enhances the defense capabilities of the FortiGate-100D.

Assumptions

As shown in other examples and network diagrams throughout this document, the Example Corporation has a pair of FortiGate-100D units in an HA cluster. To simplify this example, the cluster is replaced with a single FortiGate-100D.

An ASM-CE4 is installed in the FortiGate-100D.

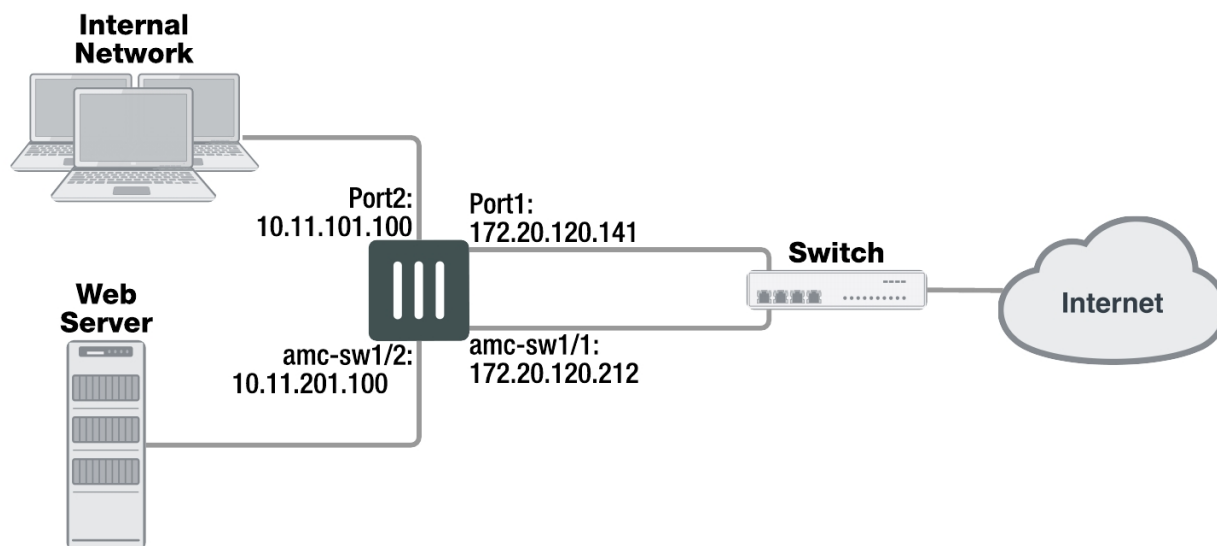
The network is configured as shown below.

Network configuration

The Example Corporation network needs minimal changes to incorporate the ASM-CE4. Interface amc-sw1/1 of the ASM-CE4 is connected to the Internet and interface amc-sw1/1 is connected to the web server.

Since the main office network is connected to port2 and the Internet is connected to port1, a switch is installed to allow both port1 and amc-sw1/1 to be connected to the Internet.

The FortiGate-100D network configuration



The switch used to connect port1 and amc-sw1/1 to the Internet must be able to handle any SYN flood, all of the legitimate traffic to the web site, and all of the traffic to and from the Example Corporation internal network. If the switch can not handle the bandwidth, or if the connection to the service provider can not provide the required bandwidth, traffic will be lost.

Security module configuration

The Fortinet security modules come configured to give equal priority to content inspection and firewall processing. The Example Corporation is using a ASM-CE4 module to defend its web server against SYN flood attacks so firewall processing is a secondary consideration.

Use these CLI commands to configure the security module in ASM slot 1 to devote more resources to content processing, including DoS and IPS, than to firewall processing.

```
config system amc-slot
  edit sw1
    set optimization-mode fw-ips
    set ips-weight balanced
    set ips-p2p disable
    set ips-fail-open enable
    set fp-disable none
    set ipsec-inb-optimization enable
    set syn-proxy-client-timer 3
    set syn-proxy-server-timer 3
  end
```

These settings do not disable firewall processing. Rather, when the security module nears its processing capacity, it will chose to service content inspection over firewall processing.

Anti-spam filter

This section describes how to configure FortiGate email filtering for IMAP, POP3, and SMTP email. Email filtering includes both spam filtering and filtering for any words or files you want to disallow in email messages. If your FortiGate unit supports SSL content scanning and inspection, you can also configure spam filtering for IMAPS, POP3S, and SMTPS email traffic.

The Anti-Spam security profile is only available when operating the FortiGate in proxy-based inspection.

The following topics are included in this section:

Anti-spam concepts

You can configure the FortiGate unit to manage unsolicited commercial email by detecting and identifying spam messages from known or suspected spam servers.

The FortiGuard Anti-Spam service uses both a sender IP reputation database and a spam signature database, along with sophisticated spam filtering tools, to detect and block a wide range of spam messages. Using FortiGuard Anti-Spam profile settings, you can opt to filter with IP address checking, URL checking, email checksum checking, detection of phishing URLs in email, and spam submission. Updates to the IP reputation and spam signature databases are provided continuously via the global FortiGuard Distribution Network.

At the [FortiGuard Anti-Spam](#) service page on the [FortiGuard Labs](#) website, you can find out whether an IP address is blacklisted in the FortiGuard Anti-Spam IP reputation database, or whether a URL or email address is in the signature database.

Anti-spam techniques

The FortiGate unit has a number of techniques available to help detect spam. Some use the FortiGuard Anti-Spam service and require a subscription. The remainder use your DNS servers or use lists that you must maintain.

Black white list

These are the types of black white lists available. They include:

- **IP/Netmask**

The FortiGate unit compares the IP address of the client delivering the email to the addresses in the IP address black / white list specified in the email filter profile. If a match is found, the FortiGate unit will take the action configured for the matching black / white list entry against all delivered email.

The default setting of the `smtp-spamhdrop` CLI command is `disable`. If enabled, the FortiGate unit will check all the IP addresses in the header of SMTP email against the specified IP address black / white list.

- **Email Wildcard**

The FortiGate unit compares the sender email address, as shown in the message header and envelope MAIL FROM, to the pattern in the patterned field. The wildcard symbol is used in the place of characters in the address that may vary from the pattern. If a match is found, the FortiGate unit will take the action configured for the matching black / white list entry.

- **Email Regular Expression**

The FortiGate unit compares the sender email address, as shown in the message envelope MAIL FROM, to the

pattern in the patterned field. The regular expression that can be used is much more sophisticated than a simple wildcard variable. If a match is found, the FortiGate unit will take the action configured for the matching black/white list entry.

Pattern

The pattern field is for entering the identifying information that will enable the filter to correctly identify the email messages.

- If the type is IP/Netmask the filter will be an IP address with a subnet mask.
- If the type is Email Wildcard the filter will be an email address with a wildcard symbol in place of the variable characters. For example *.example.com or fred@*.com.
- If the type is Email Regular Expression, regular expression can be used to create a more granular filter for email addresses. For example, `^[_a-z0-9-]+(\.[_a-z0-9-]+)*@(example|xample|examp).(com|org|net)` could be used filter based on a number of combinations of email domain names.

Action

- **Tag**
If this is the selected action, the email will be allowed through but it will be tagged with an indicator that clearly marks the email as spam.
- **Pass**
If this is the selected action, the email will be allowed to go through to its destination on the assumption that the message is not spam.
- **Discard**
If this is the selected action, the email will be dropped at the before reaching its destination.

Status

Indicates whether this particular list is enabled or disabled.

Banned word check

When you enable banned word checking, your FortiGate unit will examine the email message for words appearing in the banned word list specified in the Anti-Spam profile. If the total score of the banned word discovered in the email message exceeds the threshold value set in the Anti-Spam profile, your FortiGate unit will treat the message as spam.

When determining the banned word score total for an email message, each banned word score is added once no matter how many times the word appears in the message. Use the command `config spamfilter bword` to add an email banned word list. Use the command `config spamfilter profile` to add a banned word list to an Anti-Spam profile.

How content is evaluated

Every time the banned word filter detects a pattern in an email message, it adds the pattern score to the sum of scores for the message. You set this score when you create a new pattern to block content. The score can be any number from zero to 99999. Higher scores indicate more offensive content. When the total score equals or exceeds the threshold, the email message is considered as spam and treated according to the spam action configured in the email filter profile. The score for each pattern is counted only once, even if that pattern appears many times in the email message. The default score for banned word patterns is 10 and the default threshold is 10. This means that by default, an email message is blocked by a single match.

A pattern can be part of a word, a whole word, or a phrase. Multiple words entered as a pattern are treated as a phrase. The phrase must appear as entered to match. You can also use wildcards or regular expressions to have a pattern match multiple words or phrases.

For example, the FortiGate unit scans an email message that contains only this sentence: “The score for each word or phrase is counted only once, even if that word or phrase appears many times in the email message.”

Banned word pattern	Pattern type	Assigned score	Score added to the sum for the entire page	Comment
word	Wildcard	20	20	The pattern appears twice but multiple occurrences are only counted once.
word phrase	Wildcard	20	0	Although each word in the phrase appears in the message, the words do not appear together as they do in the pattern. There are no matches.
word*phrase	Wildcard	20	20	The wildcard represents any number of any character. A match occurs as long as “word” appears before “phrase” regardless of what is in between them.
mail*age	Wildcard	20	20	Since the wildcard character can represent any characters, this pattern is a match because “email message” appears in the message.

In this example, the message is treated as spam if the banned word threshold is set to 60 or less.

Adding words to a banned word list

When you enter a word, set the `Pattern-type` to wildcards or regular expressions.

Wildcard uses an asterisk (“*”) to match any number of any character. For example, `re*` will match all words starting with “re”.

Regular expression uses Perl regular expression syntax. See <http://perldoc.perl.org/perlretut.html> for detailed information about using Perl regular expressions.

DNS-based Blackhole List (DNSBL)

A DNSBL is a list of IP addresses, usually maintained by a third party, which are identified as being associated with spamming.

FortiGuard Anti-spam Service.

FortiGuard IP address check

The FortiGate unit queries the FortiGuard Anti-Spam Service to determine if the IP address of the client delivering the email is blacklisted. A match will cause the FortiGate unit to treat delivered messages as spam.

The default setting of the `smtp-spamhdrop` CLI command is `disable`. When you enable FortiGuard IP address checking, your FortiGate unit will submit the IP address of the client to the FortiGuard service for

checking. If the IP address exists in the FortiGuard IP address black list, your FortiGate unit will treat the message as spam.

FortiGuard URL check

When you enable FortiGuard URL checking, your FortiGate unit will submit all URLs appearing in the email message body to the FortiGuard service for checking. If a URL exists in the FortiGuard URL black list, your FortiGate unit will treat the message as spam.

FortiGuard email checksum check

When you enable FortiGuard email checksum checking, your FortiGate unit will submit a checksum of each email message to the FortiGuard service for checking. If a checksum exists in the FortiGuard checksum black list, your FortiGate unit will treat the message as spam.

Detect phishing URLs in email

When you enable FortiGuard phishing URL detection, your FortiGate unit will submit all URL hyperlinks appearing in the email message body to the FortiGuard service for checking. If a URL exists in the FortiGuard URL phishing list, your FortiGate unit will remove the hyperlink from the message. The URL will remain in place, but it will no longer be a selectable hyperlink.

FortiGuard spam submission

Spam submission is a way you can inform the FortiGuard Anti-Spam service of non-spam messages incorrectly marked as spam. When you enable this setting, the FortiGate unit adds a link to the end of every message marked as spam. You then select this link to inform the FortiGuard Anti-Spam service when a message is incorrectly marked.

Trusted IP addresses

A list of IP addresses that are trusted by the FortiGate is created. Any email traffic coming in from these IP address will be exempted to perform IP based check, such as DNSBL/RBL, FortiShield SPAM IP or locally defined IP black list check.

If the FortiGate unit sits behind a company's Mail Transfer Units, it may be unnecessary to check email IP addresses because they are internal and trusted. The only IP addresses that need to be checked are those from outside of the company. In some cases, external IP addresses may be added to the list if it is known that they are not sources of spam.

MIME header

This feature filters by the MIME header. MIME header settings are configured in a separate part of the command tree but MIME header filtering is enabled within each profile.

HELO DNS lookup

Whenever a client opens an SMTP session with a server, the client sends a HELO command with the client domain name. The FortiGate unit takes the domain name specified by the client in the HELO and does a DNS lookup to determine if the domain exists. If the lookup fails, the FortiGate unit determines that any messages delivered during the SMTP session are spam.

The HELO DNS lookup is available only for SMTP traffic.

Return email DNS check

The FortiGate unit performs a DNS lookup on the If no such record exists, the message is treated as spam.

When you enable return email DNS checking, your FortiGate unit will take the domain in the reply-to email address and reply-to domain and check the DNS servers to see if there is an A or MX record for the domain. If the domain does not exist, your FortiGate unit will treat the message as spam.

Configuring Anti-spam

FortiGuard email filtering techniques us FortiGuard services to detect the presence of spam among your email. A FortiGuard subscription is required to use the FortiGuard email filters. To enable email filtering an email filter needs to be created and then the filter needs to be associated with a security policy.

The Anti-Spam security profile is only available when operating the FortiGate in proxy-based inspection.

The filter can be created as follows:

- Go to **Security Profiles > Anti-Spam**.
 - Select the **Create New** icon (a plus symbol in a circle in the upper right hand corner).
 - Select the **List** icon (a page symbol in the upper right hand corner) and in the new window select **Create New**.

An existing filter can be edited as follows:

- Go to **Security Profiles > Anti-Spam**.
 - Select the filter that you wish to edit from the dropdown menu in the upper right corner.
 - Select the List icon (a page symbol in the upper right hand corner) and select the filter that you wish to edit from the list.

Once you are in the proper **Edit Anti-Spam Profile** window, you can enter a name in the Name field if it's a new filter.

The Comments field is for a description or other information that will assist in understanding the function or purpose of the this particular filter.

Before any of the other features or options of the filter appear the checkbox next to Enable Spam Detection and Filtering must be checked.

Spam detection by protocol

This matrix includes three rows that represent the email protocols IMAP, POP3 and SMTP.

There are also columns for:

Spam Action

For the client protocols, IMAP and POP3 the options are:

- **Tag** - This action will insert a tag into the email somewhere so that when the recipients view the email they will be warned that it is likely a spam.
- **Pass** - This action will allow any emails marked as spam to pass through without change. If this option is chosen, the Tag comments will be greyed out.

For the transfer protocol, SMTP, the options are:

- **Tag** - This action will insert a tag into the email somewhere so that when the recipients view the email they will be warned that it is likely a spam.
- **Discard** - The action will drop the email before it reaches its destination.
- **Pass** - This action will allow any emails marked as spam to pass through without change. If this option is chosen, the Tag comments will be greyed out.

Tag Location

- **Subject** - The contents of the Tag Format will be inserted into the subject line. The subject line is the most commonly used.
- **MIME** - The contents of the Tag Format will be inserted in with the MIME header header.

Tag Format

The contents of this field will be entered into the tag location specified. The most common tag is something along the lines of [Spam] or **SPAM**

FortiGuard spam filtering

The options in the section are ones that require a FortiGuard subscription.

The options available in this section, to be selected by checkbox are:

- IP Address Check
- URL Check
- Detect Phishing URLs in Email
- Email Checksum Check
- Spam Submission

Local spam filtering

The options in the section are ones can be managed on the local device without the need for a FortiGuard subscription.

The options available in this section, to be selected by checkbox are:

- HELO DNS Lookup
- Return Email DNS Check
- Black White List - checking this option will produce a table that can be edited to create a number of black / white lists that can be separately configured and enabled.

Another local spam filter profile option that can only be configured in the CLI is the `bannedword.check`. To configure this, enter the following commands in the CLI:

```
config spamfilter profile
  edit <filter_name>
    set options bannedword
    set spam-bword-table 1
  next
end
```

See the section on [banned word checking](#) for more information on how content is evaluated.

Order of spam filtering

The FortiGate unit checks for spam using various filtering techniques. The order in which the FortiGate unit uses these filters depends on the mail protocol used.

Filters requiring a query to a server and a reply (FortiGuard Anti-Spam service and DNSBL/ORDBL) are run simultaneously. To avoid delays, queries are sent while other filters are running. The first reply to trigger a spam action takes effect as soon as the reply is received.

Each spam filter passes the email to the next if no matches or problems are found. If the action in the filter is **Mark as Spam**, the FortiGate unit tags the email as spam according to the settings in the email filter profile.

For SMTP and SMTPS, if the action is **Discard**, the email message is discarded or dropped.

If the action in the filter is **Mark as Clear**, the email is exempt from any remaining filters. If the action in the filter is **Mark as Reject**, the email session is dropped.

Order of SMTP and SMTPS spam filtering

The FortiGate unit scans SMTP and SMTPS email for spam in a specific order, depending on whether or not the local override feature has been enabled. By default, local override is disabled on the FortiGate. Enabling local override will give priority to local spam filters.

You can enable local override with the CLI command `set local-override {enable | disable}` when configuring a spamfilter profile. Enable this command to override SMTP or SMTPS remote check, which includes IP RBL check, IP FortiGuard AntiSpam check and HELO DNS check, with the locally defined black/white antispam list.

SMTPS spam filtering is available on FortiGate units that support SSL content scanning and inspection.

Enabling local override of Anti-Spam filter

CLI Syntax

```
config spamfilter profile
  edit <filter_name>
    set spam-filtering enable
    set options spambwl spamfsip spamfsurl spamhelodns spamfsphish
    config smtp
      set local-override enable
    end
    set spam-bwl-table 1
  next
end
```

Order of SMTP and SMTPS spam filtering with local-override disabled

1. HELO DNS Lookup, Last Hop IP check against ORDBL
2. Return email DNS check, FortiGuard email checksum check, FortiGuard URL check, FortiGuard IP address check, Phishing URLs detection
3. Last Hop IP check local black/white list (BWL)
4. Envelope Address check local BWL
5. Headers IPs local BWL

6. Headers email address local BWL, MIME header checks based on local list of patterns (mheader)
7. Banned words (subject first, then body) based on local BWL (bword)

Order of SMTP and SMTPS spam filtering with local-override enabled

1. Last Hop IP check local black/white list (BWL)
2. Envelope Address check local BWL
3. Headers IPs local BWL, MIME header checks based on local list of patterns (mheader)
4. Headers email address local BWL
5. Banned words (subject first, then body) based on local list of patterns (bword)
6. HELO DNS Lookup, Last Hop IP check against ORDBL
7. Return email DNS check, FortiGuard email checksum check, FortiGuard URL check, FortiGuard IP address checks, Phishing URLs detection

Order of IMAP, POP3, IMAPS and POP3S spam filtering

The FortiGate unit scans IMAP, POP3, IMAPS and POP3S email for spam in the order given below. IMAPS and POP3S spam filtering is available on FortiGate units that support SSL content scanning and inspection.

1. MIME headers check, E-mail address BWL check
2. Banned word check on email subject
3. IP BWL check
4. Banned word check on email body
5. Return email DNS check, FortiGuard Antispam email checksum check, FortiGuard Antispam URL check, DNSBL & ORDBL check.

Spam actions

When spam is detected, the FortiGate unit will deal with it according to the **Spam Action** setting in the anti-spam profile. Note that POP3S, IMAPS and SMTPS spam filtering is available only on FortiGate units that support SSL content scanning and inspection. POP3, IMAP, POP3S and IMAPS mail can only be tagged. SMTP and SMTPS mail can be set to **Discard** or **Tagged**:

Discard

When the spam action is set to **Discard**, messages detected as spam are deleted. No notification is sent to the sender or recipient.

Pass

When the spam action is set to **Pass**, the spam filter is disabled for the related protocol.

Tag

When the spam action is set to **Tag**, messages detected as spam are labeled and delivered normally. The text used for the label is set in the **Tag Format** field and the label is placed in the subject or the message header, as set with the **Tag Location** option.

Anti-spam examples

Configuring simple Anti-spam protection

Small offices, whether they are small companies, home offices, or satellite offices, often have very simple needs. This example details how to enable Anti-Spam protection on a FortiGate unit located in a satellite office.

Creating an email filter profile

Most Anti-Spam settings are configured in an Anti-Spam profile. Anti-Spam profiles are selected in firewall policies. This way, you can create multiple Anti-Spam profiles, and tailor them to the traffic controlled by the security policy in which they are selected. In this example, you will create one Anti-Spam profile.

To create an Anti-Spam profile — web-based manager

1. Go to **Security Profiles > Anti-Spam**.
2. Select the **Create New** icon in the Edit Anti-Spam Profile window title.
3. In the **Name** field, enter `basic_anti-spam`.
4. Select **Enable Spam Detection and Filtering**.
5. Ensure that **IMAP**, **POP3**, and **SMTP** are selected in the header row.
These header row selections enable or disable examination of each Anti-Spam type. When disabled, the email traffic of that type is ignored by the FortiGate unit and no Anti-Spam options are available.
6. Under **FortiGuard Spam Filtering**, enable **IP Address Check**.
7. Under **FortiGuard Spam Filtering**, enable **URL Check**.
8. Under **FortiGuard Spam Filtering**, enable **E-mail Checksum Check**.
9. Select **OK** to save the email filter profile.

To create an Anti-spam profile — CLI

```
config spamfilter profile
  edit basic_anti-spam
    set options spamfsip spamfsurl spamfschksum
  end
```

Selecting the Anti-spam profile in a security policy

An Anti-Spam profile directs the FortiGate unit to scan network traffic only when it is selected in a security policy. When an Anti-Spam profile is selected in a security policy, its settings are applied to all the traffic the security policy handles.

To select the Anti-Spam profile in a security policy — web-based manager

1. Go to **Policy & Objects > IPv4 Policy**.
2. Create a new or edit a policy.
3. Turn on Anti-Spam.
4. Select the `basic_anti-spam` profile from the list.
5. Select **OK** to save the security policy.

To select the Anti-spam profile in a security policy — CLI

```
config firewall policy
  edit 1
    set utm-status enable
    set profile-protocol-options default
    set spamfilter-profile basic_anti-spam
  end
```

IMAP, POP3, and SMTP email traffic handled by the security policy you modified will be scanned for spam. Spam messages have the text “Spam” added to their subject lines. A small office may have only one security policy configured. If you have multiple policies, consider enabling spam scanning for all of them.

Blocking email from a user

Employees of the Example.com corporation have been receiving unwanted email messages from a former client at a company called example.net. The client’s email address is client@example.net. All ties between the company and the client have been severed, but the messages continue. The FortiGate unit can be configured to prevent these messages from being delivered.

To enable Anti-Spam

1. Go to **Security Profiles > Anti-Spam**.
2. Select the Anti-Spam profile that is used by the firewall policies handling email traffic from the Anti-Spam profile drop down list.
3. In the row **Tag Location**, select **Subject** for all three mail protocols.
4. In the row **Tag Format**, enter `SPAM:` in all three fields.
This means that normal spam will be tagged in the subject line.
5. Select **Enable Spam Detection and Filtering**.
6. Under **Local Spam Filtering**, enable **Black White List** and select **Create New**.
7. In the Black White List widget, select **Create New**.
8. Select **Email Address Wildcard**.
9. Enter `client@example.net` in the **Pattern** field.
 - If you wanted to prevent everyone’s email from the client’s company from getting through you could have used `*@example.net` instead.
10. Set the **Action** as **Mark as Spam**.
11. Set the **Status** to **Enable**.
12. Confirm that the SMTP protocol action is set to **Discard**.
13. Select **OK**.

Now that the email address list is created, you must enable the email filter in the Anti-Spam profile.

When this Anti-Spam profile is selected in a security policy, the FortiGate unit will reject any email message from an address ending with `@example.net` for all email traffic handled by the security policy.

Data leak prevention

The FortiGate data leak prevention (DLP) system allows you to prevent sensitive data from leaving your network. When you define sensitive data patterns, data matching these patterns will be blocked, or logged and allowed, when passing through the FortiGate unit. You configure the DLP system by creating individual filters based on file type, file size, a regular expression, an advanced rule, or a compound rule, in a DLP sensor and assign the sensor to a security policy.

Although the primary use of the DLP feature is to stop sensitive data from leaving your network, it can also be used to prevent unwanted data from entering your network and to archive some or all of the content passing through the FortiGate unit.

This section describes how to configure the DLP settings. DLP can only be configured for FortiGate units in proxy-based inspection.

The following topics are included:

- [Data leak prevention concepts](#)
- [Enable data leak prevention](#)
- [Creating or editing a DLP sensor](#)
- [DLP archiving](#)
- [DLP examples](#)

Data leak prevention concepts

Data leak prevention examines network traffic for data patterns you define through the use of the GUI and CLI commands. The DLP feature is broken down into a number of parts. Note, DLP is not available in flow-based inspection.

DLP sensor

A DLP sensor is a package of filters. To use DLP, you must enable it in a security policy and select the DLP sensor to use. The traffic controlled by the security policy will be searched for the patterns defined in the filters contained in the DLP sensor. Matching traffic will be passed or blocked according to how you configured the filters.

DLP filter

Each DLP sensor has one or more filters configured within it. Filters can examine traffic for known files using DLP fingerprints, for files of a particular type or name, for files larger than a specified size, for data matching a specified regular expression, or for traffic matching an advanced rule or compound rule.

DLP filter actions

You can configure the action taken when a match is detected. The actions include:

- Allow
- Log Only
- Block
- Quarantine IP address

Log Only is enabled by default.

Allow

No action is taken even if the patterns specified in the filter are matched.

Log Only

The FortiGate unit will take no action on network traffic matching a rule with this action. The filter match is logged, however. Other matching filters in the same sensor may still operate on matching traffic.

Block

Traffic matching a filter with the block action will not be delivered. The matching message or download is replaced with the data leak prevention replacement message.

Quarantine IP Address/ Source IP ban

Starting in FortiOS 5.2, the quarantine, as a place where traffic content was held in storage so it couldn't interact with the network or system, was removed. The term quarantine was kept to describe preventing selected source IPs from interacting with the network and protected systems. This source IP ban is kept in the kernel rather than in any specific application engine and can be queried by APIs. The features that can use the APIs to access and use the banned source IP addresses are antivirus, DLP, DoS and IPS. Both IPv4 and IPv6 version are included in this feature.

If the **quarantine-ip** action is used, the additional variable of expiry time will become available. This variable determines for how long the source IP address will be blocked. In the GUI it is shown as a field before minutes. In the CLI the option is called `expiry` and the duration is in the format `<###d##h##m>`. The maximum days value is 364. The maximum hour value is 23 and the maximum minute value is 59. The default is 5 minutes.



If a DLP sensor has contains a DLP filter with action set to **Allow** certain files and another DLP filter with action set to **Block** those same files, then the order of the filters within that sensor will determine which action is taken first.

Configuring using the CLI

To configure the DLP sensor to add the source IP address of the sender of a protected file to the quarantine or list of banned source IP addresses edit the DLP Filter, use these CLI commands:

```
config dlp sensor
  edit <sensor name>
    config filter
      edit <id number of filter>
        set action quarantine-ip
        set expiry 5m
      end
    end
  end
```

Preconfigured sensors

A number of preconfigured sensors are provided with your FortiGate unit. These can be edited to more closely match your needs.

Two of the preconfigured sensors with filters ready for you to enable are:

- Credit-Card - This sensor logs the traffic, both files and messages, that contain credit card numbers in the formats used by American Express, MasterCard and Visa.
- SSN-Sensor - This sensor logs the traffic, both files and messages, that contain Social Security Numbers with the exception of those that are WebEx invitation emails.



These rules affect only unencrypted traffic types. If you are using a FortiGate unit that can decrypt and examine encrypted traffic, you can enable those traffic types in these rules to extend their functionality if required.



Before using the rules, examine them closely to ensure you understand how they will affect the traffic on your network.

DLP document fingerprinting

One of the DLP techniques to detect sensitive data is fingerprinting (also called document fingerprinting). Most DLP techniques rely on you providing a characteristic of the file you want to detect, whether it's the file type, the file name, or part of the file contents. Fingerprinting is different in that you provide the file itself. The FortiGate unit then generates a checksum fingerprint and stores it. The FortiGate unit generates a fingerprint for all files detected in network traffic, and it is compared to all of the fingerprints stored in its fingerprint database. If a match is found, the configured action is taken.

The document fingerprint feature requires a FortiGate unit with internal storage.

Any type of file can be detected by DLP fingerprinting and fingerprints can be saved for each revision of your files as they are updated.

To use fingerprinting you:

- select the documents to be fingerprinted
- add fingerprinting filters to DLP sensors
- add the sensors to firewall policies that accept the traffic to which to apply fingerprinting.

Fingerprinting

Fingerprint scanning allows you to create a library of files for the FortiGate unit to examine. It will create checksum fingerprints so each file can be easily identified. Then, when files appear in network traffic, the FortiGate will generate a checksum fingerprint and compare it to those in the fingerprint database. A match triggers the configured action.

You must configure a document source or uploaded documents to the FortiGate unit for fingerprint scanning to work.

Fingerprinted documents

The FortiGate unit must have access to the documents for which it generates fingerprints.

Configuring the document source

To configure a DLP fingerprint document source in FortiOS 5.6.0, you must use CLI commands.

```
config dlp fp-doc-source
```

```
edit <name_str>
  set name <string>
  set server-type {smb}
  set server <string>
  set period {none | daily | weekly | monthly}
  set vdom {mgmt | current}
  set scan-subdirectories {enable | disable}
  set remove-deleted {enable | disable}
  set keep-modified {enable | disable}
  set username <string>
  set password <password>
  set file-path <string>
  set file-pattern <string>
  set sensitivity <string>
  set tod-hour <integer>
  set tod-min <integer>
  set weekday {sunday | monday | tuesday | wednesday | thursday | friday | saturday}
  set date <integer>
end
```

Configuring a DLP fingerprint sensor

To configure a DLP fingerprint sensor in FortiOS 5.6.0, you must use CLI commands.

```
config dlp sensor
  edit <sensor name>
    config filter
      edit <id number of filter>
        set proto {smtp | pop3 | imap http-get | http-post | ftp | nntp | mapi}
        set filter-by fingerprint
        set fp-sensitivity { critical | private | warning}
        set action {allow | log-only | block | ban | quarantine-ip | quarantine-
          port}
      next
    end
  next
```

Once you have set the document source and configured the DLP sensor for fingerprinting, add the DLP sensor to the applicable firewall policy. This can be done through the GUI.

File size

This filter-type checks for files exceeding a configured size. All files larger than the specified size are subject to the configured action. The value of the field is measured in kilobytes (KB).

DLP filtering by specific file types

File filters use file filter lists to examine network traffic for files that match either file names or file types. For example, you can create a file filter list that will find files called secret.* and also all JPEG graphic files. You can create multiple file filter lists and use them in filters in multiple DLP sensors as required.

Specify File Types is a DLP option that allows you to block files based on their file name or their type.

- **File types** are a means of filtering based on examination of the file contents, regardless of the file name. If you block the file type **Archive (zip)**, all zip archives are blocked even if they are renamed with a different file extension. The FortiGate examines the file contents to determine what type of file it is and then acts accordingly.
- **File Name patterns** are a means of filtering based purely on the names of files. They may include wildcards (*). For example, blocking *.scr will stop all files with an .scr file extension, which is commonly used for Windows screen saver files. Files trying to pass themselves off as Windows screen saver files by adopting the file-naming convention will also be stopped.
 - Files can specify the full or partial file name, the full or partial file extension, or any combination. File pattern entries are not case sensitive. For example, adding *.exe to the file pattern list also blocks any files ending with .exe.
 - Files are compared to the enabled file patterns from top to bottom, in list order.



If DLP detects a file inside an archive that should be blocked, the entire archive will be blocked.

Watermarking

Watermarking is essentially marking files with a digital pattern to mark the file as being proprietary to a specific company. Fortinet provides a Linux-based utility that applies a digital watermark to files. The utility adds a small (approx. 100 byte) pattern to the file that is recognized by the DLP watermark filter. The pattern is invisible to the end user.

When watermarking a file it should be verified that the pattern matches up to a category found on the FortiGate firewall. For example, if you are going to watermark a file with the sensitivity level of “Secret” you should verify that “Secret” is a sensitivity level that has been assigned in the FortiGate unit.

Watermark Sensitivity

If you are using watermarking on your files you can use this filter to check for watermarks that correspond to sensitivity categories that you have set up.

The Corporate Identifier is to make sure that you are only blocking watermarks that your company has place on the files, not watermarks with the same name by other companies.

Software Versions

Before planning on using watermarking software it is always best to verify that the software will work with your OS. Currently, the only utility available to watermark files is a Linux-based command line tool. It is available for download from the [Fortinet Customer Service & Support](#) website, with a valid support contract and access to the site. To access the file:

1. Sign into the [Fortinet Customer Service & Support](#) website.
2. Go to <https://support.fortinet.com/Download/FirmwareImages.aspx>.
3. Navigate to the image file path for /FortiGate / v5.00 / 5.0 / WATERMARK
4. Download the file **fortinet-watermark-linux.out**.

File types

The watermark utility does not work with every file type. The following file types are supported by the watermark tool: .txt; .pdf; .doc; .xls; .ppt; .docx; pptx; and, .xlsx.

Syntax of the watermark utility

The tool is executed in a Linux environment by passing in files or directories of files to insert a watermark.

USAGE:

```
watermark_linux_amd64 <options> -f <file name> -i <identifier> -l <sensitivity level>
watermark_linux_amd64 <options> -d <directory> -i <identifier> -l <sensitivity level>
```

Options:

```
-h print help
-I inplace watermarking (don't copy file)
-o output file (or directory in directory mode)
-e encode <to non-readable>
-i add watermark identifier
-l add watermark sensitivity level
-D delete watermark identifier
-L delete watermark sensitivity level
```

Regular expression

The FortiGate unit checks network traffic for the regular expression specified in a regular expression filter. The regular expression library used by Fortinet is a variation of a library called PCRE (Perl Compatible Regular Expressions). A number of these filters can be added to a sensor making a sort of 'dictionary' subset within the sensor.

Some other, more limited DLP implementations, use a list of words in a text file to define what words are searched for. While the format used here is slightly different than what some people are used to, the resulting effect is similar. Each regular expression filter can be thought of as a more versatile word to be searched against. In this dictionary (or sensor), the list of words is not limited to just predefined words. It can include expressions that accommodate complex variations on those words and even target phrases. Another advantage of the individual filter model of this dictionary over the list is that each word can be assigned its own action, making this implementation much more granular.

Encrypted

This filter is a binary one. If the file going through the policy is encrypted the action is triggered.

Examining specific services

To assist in optimizing the performance of the firewall, the option exists to select which services or protocol traffic will be checked for the targeted content. This setting gives you a tool to save the resources of the FortiGate unit by only using processing cycles on the relevant traffic. Just check the boxes associated with the service / protocol that you want to have checked for filter triggers.

Enable data leak prevention

DLP examines your network traffic for data patterns you specify. The FortiGate unit then performs an action based on the which pattern is found and a configuration set for each filter trigger.

DLP is not available in flow-based inspection.

General configuration steps

Follow the configuration procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. [Create a DLP sensor](#).
New DLP sensors are empty. You must create one or more filters in a sensor before it can examine network traffic.
2. Add one or more filters to the DLP sensor.
Each filter searches for a specific data pattern. When a pattern in the active DLP sensor appears in the traffic, the FortiGate unit takes the action configured in the matching filter. Because the order of filters within a sensor cannot be changed, you must configure DLP in sequence.
3. Add the DLP sensor to one or more firewall policies that control the traffic to be examined.

Creating/editing a DLP sensor

DLP sensors are collections of filters. You must also specify an action for the filter when you create it in a sensor. Once a DLP sensor is configured, you add it to a security policy profile. Any traffic handled by that security policy will be examined according to the DLP sensor configuration.

DLP is not available in flow-based inspection.

To create/edit a DLP sensor in the GUI

1. Go to **Security Profiles > Data Leak Prevention**.
2. Choose whether you want to edit an existing sensor or create a new one.
 - The default sensor is the one displayed by default.
 - To edit an existing sensor, select it by either using the drop down menu in the upper right hand corner of the window or by selecting the **List** icon (the furthest right of the 3 icons in the upper right of the window, resembling a page with some lines on it), and then selecting the profile you want to edit from the list.
 - To create a new sensor, select the **Create New** icon (a plus sign within a circle) or the List icon and then select the Create New link in the upper left corner of the window that appears.
3. Enter a name in the **Name** field for any new DLP sensors.
4. Optionally, you may also enter a comment. The comment appears in the DLP sensor list and can remind you of the details of the sensor.
5. At this point you can add filters to the sensor (see adding filters to a DLP sensor) or select **OK** to save the sensor. Without filters, the DLP sensor will do nothing.

Adding filters to a DLP sensor

Once you have created a DLP sensor, you need to add filters.

1. To add filters to a DLP sensor
2. Go to **Security Profiles > Data Leak Prevention**.
3. Select the sensor you wish to edit using the drop-down menu or the sensor list window.
4. Within the Edit DLP Sensor window select **Create New**. A New Filter window should pop up.
5. Select the type of filter. You can choose either **Messages** or **Files**, depending on which of these two are chosen different options will be available.

Message filter will have these configuration options:

- [radio button] Containing: [drop-down menu including: Credit Card # or SSN]
- [radio button] Regular Expression [input field]

Examine the following services:

Web Access

- HTTP-POST

Email

- [check box] SMTP
- [check box] POP3
- [check box] IMAP
- [check box] MAPI

Others

- [check box] NNTP

Action [from drop-down menu]

- Allow
- Log Only (default)
- Block
- Quarantine IP address

Files filter will allow you to choose one of these options:

- **Containing:** drop-down menu including: Credit Card # or SSN
- **File Size** > []KB files greater than the number of KB entered
- Specify File Types
File Types: ["Click to add..."drop-down menu of File extensions]
File Name Patterns:["Click to add..."drop-down menu]
- [radio button] Regular Expression [input field]
- [radio button] Encrypted

Examine the following Services:

Web Access

- [check box] HTTP-POST
- [check box] HTTP-GET

Email

- [check box] SMTP
- [check box] POP3
- [check box] IMAP
- [check box] MAPI

Others

- [check box] FTP
- [check box] NNTP

Action [from drop-down menu]

- Allow
 - Log Only (default)
 - Block
 - Quarantine IP address
6. Select **OK**.
 7. Repeat Steps 6 and 7 for each filter.
 8. Select **Apply** to confirm the settings of the sensor.



If you have configured DLP to block IP addresses and if the FortiGate unit receives sessions that have passed through a NAT device, all traffic from that NAT device — not just traffic from individual users — could be blocked. You can avoid this problem by implementing authentication.



To view or modify the replacement message text, go to **System > Replacement Messages**.

DLP archiving

DLP is typically used to prevent sensitive information from getting out of your company network, but it can also be used to record network use. This is called DLP archiving. The DLP engine examines email, FTP, NNTP, and web traffic. Enabling archiving for rules when you add them to sensors directs the FortiGate unit to record all occurrences of these traffic types when they are detected by the sensor.

Since the archive setting is configured for each rule in a sensor, you can have a single sensor that archives only the things you want.

You can archive Email, FTP, HTTP, and session control content:

- Email content includes IMAP, POP3, and SMTP sessions. Email content can also include email messages tagged as spam by Email filtering. If your unit supports SSL content scanning and inspection, email content can also include IMAPS, POP3S, and SMTPS sessions.
- HTTP content includes HTTP sessions. If your unit supports SSL content scanning and inspection HTTP content can also include HTTPS sessions.

DLP archiving comes in two forms: **Summary** and **Full**.

Summary archiving records information about the supported traffic types. For example, when an email message is detected, the sender, recipient, message subject, and total size are recorded. When a user accesses the Web, every URL the user visits recorded. The result is a summary of all activity the sensor detected.

For more detailed records, use full archiving. When an email message is detected, the message itself, including any attachments, is archived. When a user accesses the Web, every page the user visits is archived. Far more detailed than a summary, full DLP archives require more storage space and processing.

Because both types of DLP archiving require additional resources, DLP archives are saved to a FortiAnalyzer unit or the FortiGuard Analysis and Management Service (subscription required).

You can use DLP archiving to collect and view historical logs that have been archived to a FortiAnalyzer unit or the FortiGuard Analysis and Management Service. DLP archiving is available for FortiAnalyzer when you add a

FortiAnalyzer unit to the Fortinet configuration. The FortiGuard Analysis server becomes available when you subscribe to the FortiGuard Analysis and Management Service.

Two sample DLP sensors are provided with DLP archiving capabilities enabled. If you select the `Content_Summary` sensor in a security policy, it will save a summary DLP archive of all traffic the security policy handles. Similarly, the `Content_Archive` sensor will save a full DLP archive of all traffic handled the security policy you apply it to. These two sensors are configured to detect all traffic of the supported types and archive them. You can see these sensors in the GUI but the configuration is only visible through the CLI; DLP archiving is set in the CLI only.

To set the archive to Summary

```
config dlp sensor
  edit <name of sensor>
    set summary-proto smtp pop3 imap http ftp nntp msn yahoo mapi
  end
```

To set the archive to Full

```
config dlp sensor
  edit <name of sensor>
    set full-archive-proto smtp pop3 imap http ftp nntp msn yahoo mapi
  end
```



If you set the `full-archive-proto` filter to include one or more of the protocols set by the `proto` option, then the archive action is disabled.

DLP examples

You can configure DLP sensors and filters when your FortiGate is operating in proxy-based inspection.

- [Blocking content with credit card numbers](#)
- [Blocking emails larger than 15 MB and logging emails from 5 MB to 15 MB](#)
- [Blocking selectively based on a fingerprint](#)

Blocking content with credit card numbers

When the objective is to block credit card numbers one of the important things to remember is that two filters will need to be used in the sensor. One filter is to prevent sensitive files from being leaked and another is to retain any sensitive data that is not a file (for example, messages or email content).

In the default Credit-Card sensor, you will notice a few things.

- The **Action** is set to **Log Only**
- In the **Files** filter not all of the services are being examined.

If you wish to block as much content as possible with credit card numbers in it instead of just logging most the traffic that has it, the existing sensor will have to be edited.

1. Go to **Security Profiles > Data Leak Prevention**.

Some configurations will have a preconfigured Credit Card sensor where you can use the drop down menu to select **Credit-Card**. If your configuration doesn't already have one create a new sensor.

2. Use the **Create New** icon to add a new sensor.
3. *Create/edit the first filter.* Set **Type** to **Messages** and select **Containing Credit Card #**.
4. Go to **Examine the Following Services** and select all services .
5. Set **Action** to **Block**.
6. Select **OK** or **Apply**.
7. *Create/edit the second filter.* Set **Type** to **Files** and select **Containing Credit Card #**.
8. Go to **Examine the Following Services** and select all services .
9. Set **Action** to **Block**.
10. Select **OK** or **Apply**.
11. Edit the appropriate policies so that under **Security Profiles**, **DLP** is turned on and the **Credit-Card** sensor is selected.

Blocking emails larger than 15 MB and logging emails from 5 MB to 15 MB

Multiple filters will have to be used in this case and the order that they are used is important. Because there is no mechanism to move the filters within the sensor the order that they are added to the sensor is important.

1. Go to **Security Profiles > Data Leak Prevention**.
2. Use the **Create New** icon to add a new sensor. Give it a descriptive **Name**, such as *block_large_emails*. Optionally, enter a descriptive comment.

Once the sensor has been created, a new filter will need to be added.

3. *Create the filter to block the emails over 15 MB.* In the filters table select **Create New**.
4. Set **Type** to **Messages** and enter 15360 in the field next to **File size over**. (1MB = 1024KB, 15 MB = 15 x 1024KB = 15360KB)
5. Go to **Examine the Following Services** and select all **Email** services .
6. Set **Action** to **Block**.
7. Select **OK**.
8. *Create the filter to log emails between 5 MB and 10 MB.* In the filters table select **Create New**.
9. Set **Type** to **Files**.
10. Enter 5120 in the field next to **File size over**. (1MB = 1024KB, 5 MB = 5 x 1024KB = 5124KB)
11. Go to **Examine the Following Services** and select all the email services .
12. Set action to **Log Only**.
13. Select **OK**.

The reason that the block filter is placed first is because the filters are applied in sequence and once the traffic triggers a filter, the action is applied and then the traffic is passed on to the next test. If the Log Only filter which checks for anything over 1MB is triggered this would include traffic over 15MB, so a 16 MB file would only be logged. In the described order, the 16 MB file will be blocked and the 3 MB file will be logged.

Blocking selectively based on a fingerprint

The following is a fairly complex example but shows what can be done by combining various components in the correct configuration.

The company has a number of copyrighted documents that it does not want “escaping” to the Internet but it does want to be able to send those documents to the printers for turning into hardcopy.

The policies and procedures regarding this issue state that:

- Only members of the group **Senior_Editors** can send copyrighted material to the printers.
- Every member of the company by default is included in the group **employees**.
- Even permitted transmission of copyrighted material should be recorded.
- All of the printers IP addresses are in a group called **approved_printers**.
- There is a file share called **copyrighted** where any file that is copyrighted is required to have a copy stored.
- It doesn't happen often but for legal reasons sometimes these files can be changed, but all versions of a file in this directory need to be secured.
- All network connections to the Internet must have AntiVirus enabled using at least the default profile.
- The SSL/SSH Inspection profile used will be **default**.

It is assumed for the purposes of this example that:

- Any addresses or address groups have been created.
- User accounts and groups have been created.
- The account used by the FortiGate is fgtaccess.
- The copyrighted sensitivity level needs to be created.
- The copyrighted material is stored at \\192.168.27.50\books\copyrighted\

1. Add a new Sensitivity Level by running the following commands in the CLI:

```
config dlp fp-sensitivity
  edit copyrighted
end
```

2. Apply files to the fingerprint database by running these commands in the CLI:

```
config dlp fp-doc-source
  edit "copyrighted_material"
    set server-type smb
    set server 192.168.27.50
    set username fgtaccess
    set password *****
    set file-path books/copyrighted/
    set file-pattern *.pdf
    set sensitivity copyrighted
    set period daily
    set tod-hour 2
    set tod-min 0
    set scan-subdirectories enable
    set remove-deleted disable
    set keep-modified enable
  next
end
```

Two Sensors need to be created. One for blocking the transmission of copyrighted material and a second for allowing the passing of copyrighted material under specific circumstances.

3. Create the first DLP Sensor with the following commands in CLI:

```
config dlp sensor
  edit block_copyrighted
```

```

config filter
  edit 1
    set proto smtp pop3 imap http-get http-post ftp nntp mapi
    set filter-by fingerprint
    set fp-sensitivity copyrighted
    set action block
  next
end
next

```

4. Create the second DLP Sensor

```

config dlp sensor
  edit allow_copyrighted
    config filter
      edit 2
        set proto smtp pop3 imap http-get http-post ftp nntp mapi
        set filter-by fingerprint
        set fp-sensitivity copyrighted
        set action log-only
      next
    end
  next

```

5. Create a policy to allow transmission of copyrighted material.

- a. Go to **Policy & Objects > IPv4 Policy**.
- b. Select **Create New**.
- c. Use the following values in the policy:

Incoming Interface	LAN
Source Address	all
Outgoing Interface	wan1
Destination Address	all
Schedule	always
Service	all
Action	ACCEPT
Enable NAT	enabled -- Use Destination Interface Address
AntiVirus	<ON> default
DLP	<ON> Copyrighted
SSL/SSH Inspection	<ON> default
Enable this policy	<ON>

This policy should be placed as close to the beginning of the list of policies so the it is among the first

tested against.

6. Create a policy to block transmission of copyrighted material.

This will in effect be the default template for all following policies in that they will have to use the DLP profile that blocks the transmission of the copyrighted material.

- a. Go to **Policy & Objects > IPv4 Policy**.

- b. Select **Create New** or Edit an existing policy.

- c. Use the following values in the Policy:

The fields should include what ever values you need to accomplish your requirements are but each policy should include the DLP sensor block_copyrighted. Alternatively, if a different DLP configuration is required it should include a filter that blocks **copyrighted** fingerprinted file.

If you need to create a policy that is identity based make sure that there is an Authentication rule for the group **employees** that uses the DLP sensor that blocks copyrighted material.

ICAP support

ICAP is the acronym for Internet Content Adaptation Protocol. The purpose of the feature is to offload work that would normally take place on the firewall to a separate server specifically set up for the specialized processing of the incoming traffic. This takes some of the resource strain off of the FortiGate firewall leaving it to concentrate its resources on things that only it can do.

Offloading value-added services from Web servers to ICAP servers allows those same web servers to be scaled according to raw HTTP throughput versus having to handle these extra tasks.

ICAP servers are focused on a specific function, for example:

- Ad insertion
- Virus scanning
- Content translation
- HTTP header or URL manipulation
- Language translation
- Content filtering

The following topics are included in this section:



ICAP does not appear by default in the GUI. You must enable it in **System > Feature Visibility** .

The protocol

ICAP is an Application layer protocol; its specifications are set out in [RFC 3507](#). It is, in essence, a lightweight protocol for executing a "remote procedure call" on HTTP messages and is a member of the member of the TCP/IP suite of protocols.

The default TCP that is assigned to it is 1344. Its purpose is to support HTTP content adaptation by providing simple object-based content vectoring for HTTP services. ICAP is usually used to implement virus scanning and content filters in transparent HTTP proxy caches. Content adaptation refers to performing the particular value added service, or content manipulation, for an associated client request/response.

Essentially it allows an ICAP client, in this case the FortiGate firewall, to pass HTTP messages to an ICAP server like a remote procedure call for the purposes of some sort of transformation or other processing adaptation. Once the ICAP server has finished processing the content, the modified content is sent back to the client.

The messages going back and forth between the client and server are typically HTTP requests or HTTP responses. While ICAP is a request/response protocol similar in semantics and usage to HTTP/1.1 it is not HTTP nor does it run over HTTP, as such it cannot be treated as if it were HTTP. For instance, ICAP messages can not be forwarded by HTTP surrogates.

Offloading using ICAP

If you enable ICAP in a security policy, HTTP traffic intercepted by the policy is transferred to an ICAP server in the ICAP profile added to the policy. Responses from the ICAP server are returned to the FortiGate unit which forwards them to an HTTP client or server.

You can offload HTTP responses or HTTP requests (or both) to the same or different ICAP servers.

If the FortiGate unit supports HTTPS inspection, HTTPS traffic intercepted by a policy that includes an ICAP profile is also offloaded to the ICAP server in the same way as HTTP traffic.

When configuring ICAP on the FortiGate unit, you must configure an ICAP profile that contains the ICAP server information; this profile is then applied to a security policy.

Configuring ICAP

You will need to configure an ICAP [server](#) and an ICAP [profile](#).

ICAP servers

1. Go to **Security Profiles > ICAP Servers** and click on **Create New**.
2. Enter a **Name** for the server.
3. Enter the server's **IP Address**. Depending on whether you've set the IP version to 4 or 6 will determine the format that the content of this field will be set into. In the GUI it looks like the same field with a different format but in the CLI it is actually 2 different fields named "ip-address" and ip6-address.
4. Set the **Port**; 1344 is default TCP port used for the ICAP traffic. The range can be from 1 to 65535.

Maximum Connections

This value refers to the maximum number of concurrent connections that can be made to the ICAP server. The default setting is 100. This setting can only be configured in the CLI.

The syntax is:

```
config icap server
  edit <icap_server_name>
    set max-connections <integer>
  end
```

Profiles

1. Go to **Security Profiles > ICAP** and click on **Create New**.
2. Enter a **Name** for the server.
3. Enable settings as required.
 - a. **Enable Request Processing** allows the ICAP server to process request messages. If enabled this setting will also require:
 - **Server** - This is the name of the ICAP server. It is chosen from the drop down menu in the field. The servers are configure in the Security Profiles > ICAP > Server section.
 - **Path** - This is the path on the server to the processing content. For instance if the Windows share name was "Processes" and the directory within the share was "Content-Filter" the path would be "/Processes/Content-Filter/"
 - **On Failure** - There are 2 options: **Error** or **Bypass**.

- b. **Enable Response Processing** allows the ICAP server to process response messages. If enabled this setting will also require:
 - **Server** - This is the name of the ICAP server. It is chosen from the drop down menu in the field. The servers are configured in the Security Profiles > ICAP > Server section.
 - **Path** - This is the path on the server to the processing component. For instance if the Windows share name was "Processes" and the directory within the share was "Content-Filter" the path would be "/Processes/Content-Filter/".
 - **On Failure** - There are 2 options. You can choose by the use of radio buttons either **Error** or **Bypass**.
- c. **Enable Streaming Media Bypass** allows streaming media to ignore offloading to the ICAP server.
- 4. Select **Apply**.

Example ICAP sequence

This example is for an ICAP server performing web URL filtering on HTTP requests

1. A user opens a web browser and sends an HTTP request to connect to a web server.
2. The FortiGate unit intercepts the HTTP request and forwards it to an ICAP server.
3. The ICAP server receives the request and determines if the request is for URL that should be blocked or allowed.
 - If the URL should be blocked the ICAP server sends a response to the FortiGate unit. The FortiGate unit returns this response to the user's web browser. This response could be a message informing the user that their request was blocked.
 - If the URL should be allowed the ICAP server sends a request to the FortiGate unit. The FortiGate unit forwards the request to the web server that the user originally attempted to connect to.
 - When configuring ICAP on the FortiGate unit, you must configure an ICAP profile that contains the ICAP server information; this profile is then applied to a security policy.

Example ICAP scenario

Information relevant to the following example:

- The ICAP server is designed to do proprietary content filtering specific to the organization so it will have to receive the messages and send back appropriate responses.
 - The content filter is a required security precaution so if the message cannot be processed it is not allowed through.
 - Resources on both the FortiGate and the ICAP server are considerable so the maximum connections setting will be set at a double the default value to analyze the impact on performance.
 - The ICAP server's IP address is 172.16.100.55.
 - The path to the processing component is "/proprietary_code/content-filter/".
 - Streaming media is not something that the filter considers, but is allowed through the policy so processing it would be a waste of resources.
 - The ICAP profile is to be added to an existing firewall policy.
 - It is assumed that the display of the policies has already been configured to show the column "ID".
1. Enter the following to configure the ICAP server:

Go to **Security Profiles > ICAP Servers**.

Use the following values:

Name	content-filtration-server4
IP Type	IPv4
IP Address	172.16.100.55
Port	1344

Use the CLI to set the max-connections value.

```
config icap server
  edit content-filtration-server4
    set max-connections 200
  end
```

2. Enter the following to configure the ICAP profile to then apply to a security policy:

Use the following values:

Name	Prop-Content-Filtration
Enable Request Processing	enable
Server	content-filtration-server4
Path	/proprietary_code/content-filter/
On Failure	Error
Enable Response Processing	enable
Server	content-filtration-server4
Path	/proprietary_code/content-filter/
On Failure	Error
Enable Streaming Media Bypass	enable

3. Apply the ICAP profile to policy:

The purposes of this particular ICAP profile is to filter the content of the traffic coming through the firewall via policy ID#17.

- a. Go to **Policy & Objects > IPv4 Policy**.
- b. Open the existing policy ID# 17 for editing.
- c. Go to the section **Security Profiles**.
- d. Select the button next to **ICAP** so that it indicates that it's status is **ON**.
- e. Select the field with the profile name and use the drop down menu to select **Prop-Content-Filtration**.
- f. Select **OK**.

FortiClient Compliance Profiles

This section describes the FortiClient Compliance Profiles endpoint protection features and configuration.

FortiClient Compliance Profiles are used primarily to make sure connected devices are compliant with Endpoint Control and to protect against vulnerabilities. Both **Endpoint Vulnerability Scan on Client** and **System compliance** are enabled by default, while other settings are disabled by default. This allows FortiClient to work as part of a Security Fabric.



FortiClient Profiles was renamed **FortiClient Compliance Profiles** to clarify that this profile only creates "compliance rules" and cannot be used to "provision FortiClient endpoints".

You must first enable this feature. Go to System > Feature Visibility and enable Endpoint Control. This will reveal the Security Profiles > FortiClient Compliance menu item.

The following topics are included in this section:

Endpoint protection overview

Endpoint Protection enforces the use of up-to-date FortiClient Endpoint Security software on endpoints (workstation computers and mobile devices). It pushes a FortiClient profile to the FortiClient application, specifying security settings, including:

- Real-time antivirus protection - on or off
- FortiClient web category filtering based on web filters defined in a FortiGate Web Filter profile
- FortiClient Application Control (application firewall) using application sensors defined in the FortiGate Application Control profile

The FortiClient profile can also:

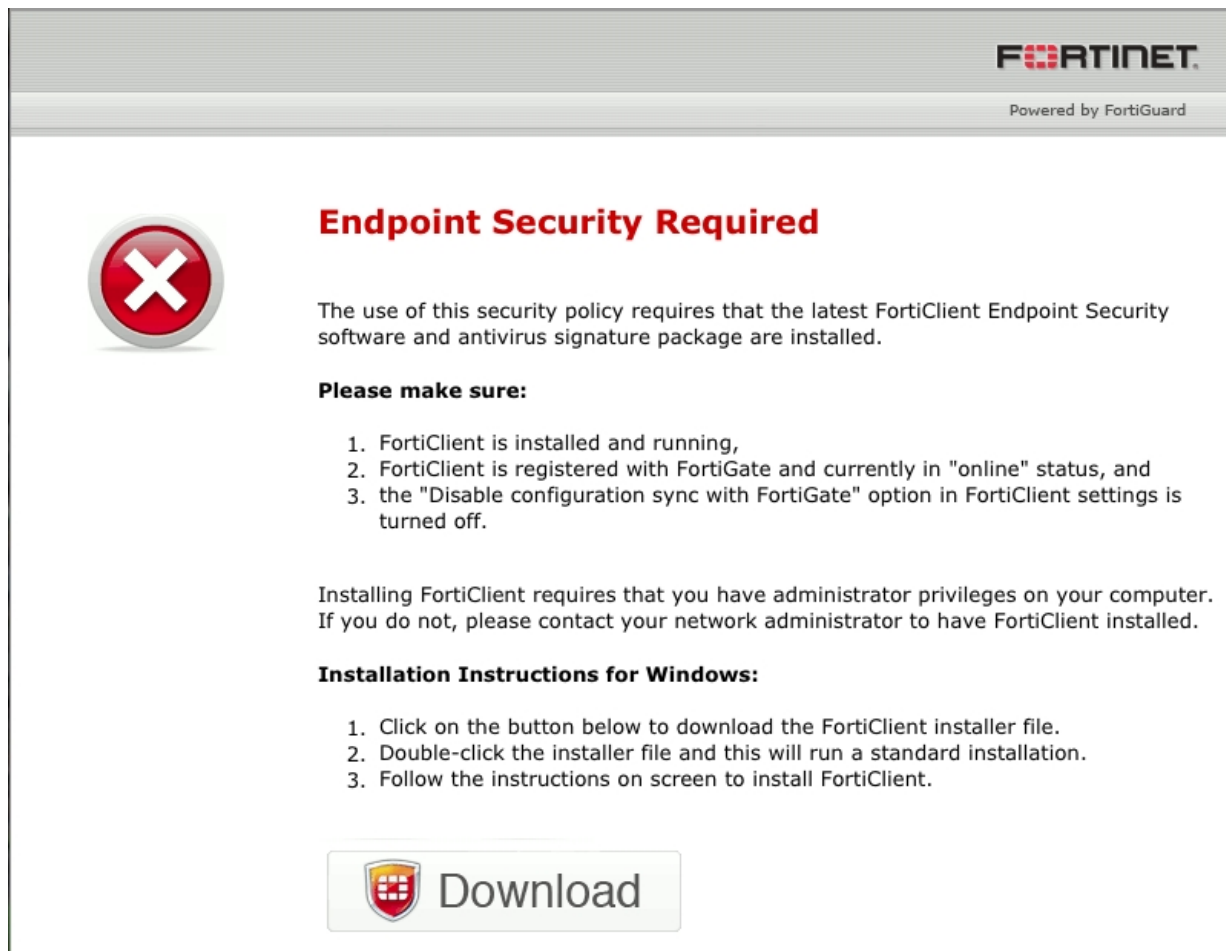
- Create VPN configurations
- Install CA certificates
- Upload logs to FortiAnalyzer or FortiManager
- Enable use of FortiManager for client software/signature update
- Enable a dashboard banner
- Enable client-based logging while on-net
- Output a mobile configuration profile (.mobileconfig file for iOS)

User experience

When using a web browser, the user of a non-compliant endpoint receives a replacement message HTML page from the FortiGate unit. The message explains that the user needs to install FortiClient Endpoint Security and provides a link to do so. The user cannot continue until the FortiClient software is installed.

For information about modifying the replacement message, see [Modifying the endpoint protection replacement messages on page 2552](#).

Default FortiClient non-compliance message for Windows



After installing FortiClient Endpoint Security, you will receive an invitation to register with the FortiGate unit. If you accept the invitation, the FortiClient profile is sent to the device's FortiClient application. Now the device is compliant and can connect to the network. FortiClient Endpoint Security registered with a FortiGate unit does not need to be separately licensed with FortiGuard.

The FortiGate unit can also register endpoints connecting over the Internet through a VPN. See [Configuring endpoint registration over a VPN on page 2548](#).

Licensing and FortiGate endpoint registration limits

To view the number of endpoints that are registered and the total that can be registered, go to **Dashboard**. Under **Licenses**, find **FortiClient**. You will see text like "4 / 10". This means that there are four registered endpoints and a total of ten are allowed.

When the registration limit is reached, the next FortiClient-compatible device will not be able to register with the FortiGate unit. A message appears in the FortiClient application. The FortiClient profile is not sent to client and the client cannot connect through the FortiGate unit.

For all FortiGate models, the maximum number of registered endpoints is ten. For all models except 20C, you can purchase an endpoint license to increase this capacity:

To add an endpoint license - GUI

1. Go to **Dashboard**.
2. In the **Licenses** widget, click on **FortiClient**, select **Enter License**.
3. Enter the license key in the window that slides in from the right, and select **OK**.

Maximum registered endpoints with endpoint license

FortiClient endpoint licenses for FortiOS 5.6.0 can be purchased in multiples of 100. There is a maximum client limit based on the FortiGate's model. FortiCare enforces the maximum limits when the customer is applying the license to a model.

If you are using the ten free licenses for FortiClient, support is provided on the Fortinet Forum (forum.fortinet.com). Phone support is only available for paid licenses.

Model(s)	Maximum Client Limit
VM00	200
FGT/FWF 30 to 90 series	200
FGT 100 to 400 series	600
FGT 500 to 900 series, VM01, VM02	2,000
FGT 1000 to 2900 series	20,000
FGT 3000 to 3600 series, VM04	50,000
FGT 3700D and above, VM08 and above	100,000

Older FortiClient SKUs will still be valid and can be applied to FortiOS 5.4 and 5.6.

Configuring endpoint protection

Endpoint Protection requires that all hosts connecting to an interface have the FortiClient Endpoint Security application installed. Make sure that all endpoints behind the interface are able to install this application. Currently, FortiClient Endpoint Security is available for Microsoft Windows (2000 and later), Apple (Mac OS X and later), and Android devices only.

By default, the FortiGuard service provides the FortiClient installer. If you prefer to host it on your own server, see [Changing the FortiClient installer download location](#), below.

To set up Endpoint Protection, complete the following:

- Create a FortiClient Profile or use the default profile. See [Creating a FortiClient profile on page 2544](#). Enable the application sensor and web category filtering profiles that you want to use.
- Configure the FortiGate unit to support endpoint registration using FortiTelemetry (under **Network > Interfaces**, allow FortiTelemetry admission control).

- Optionally, enforce FortiClient registration. See [Enforcing FortiClient registration on page 2545](#).
- Optionally, configure application sensors and web filter profiles as needed to monitor or block applications.
- Optionally, modify the **Endpoint NAC Download Portal** replacement messages (one per platform). See [Modifying the endpoint protection replacement messages on page 2552](#).

Creating a FortiClient profile

FortiClient profiles allow you to perform vulnerability scans on endpoints and make sure endpoints are running compliant versions of FortiClient. Also, security posture features cause FortiClient to apply realtime protection, AntiVirus, web filtering, and application control on endpoints.

It is possible for more than one profile to be assigned to a device type. As with security policies, clients are matched to FortiClient profiles in the order that the profiles appear in the list.

Features involving general settings have been removed from the FortiClient profile GUI in 5.4.1. Features emphasizing compliance of the endpoint devices have been added. These enhancements facilitate integration with the Security Fabric.

To create a FortiClient profile - GUI

1. If you plan to use the Application Firewall feature in the FortiClient profile, go to **Security Profiles > Application Control** to create the Application Sensors that you will need.
2. If you plan to use the Web Category Filtering, go to **Security Profiles > Web Filter** to create the Web Filter Profile that you will need.
3. Go to **Security Profiles > FortiClient Compliance**. If there is only the default FortiClient profile, it will be displayed and ready to edit. At the top right of the page you can select or create other profiles.
4. Select **Create New** or edit an existing profile.
5. In **Assign Profile To**, select the device groups, user groups, and users to which this FortiClient profile applies. ***This is not available for the default profile.***
6. Set the **Endpoint Vulnerability Scan on Client** quarantine level. Similar to FortiOS 5.4, you can set the FortiClient Profile to run the FortiClient vulnerability scanner on endpoints and you can set the vulnerability quarantine level to quarantine endpoints that don't comply. The FortiGate will quarantine a host when a vulnerability with the level of severity selected, or higher, is detected. Options are: **Critical, High, Medium, Low, and Information**.
7. **System Compliance** FortiOS 5.6 system compliance settings are similar to those in 5.4 with the addition of a non-compliance action. System compliance checking is performed by FortiClient but the non-compliance action is applied by the FortiGate:
 - select the **Minimum FortiClient version**, if necessary. The lowest supported version is 5.4.1.
 - identify which logs, if any, you will upload to FortiAnalyzer
 - set the **Non-compliance action: Block or Warning**.
8. Under **Security Posture Check**, enable the required options for your network:
 - **Realtime Protection**
 - **Third party AntiVirus on Windows** is required for Windows endpoints
 - identify which logs, if any, you will upload to FortiAnalyzer
 - select whether to enable an **Web Filter** security profile, and / or an **Application Control** sensor.
 - set the **Non-compliance action: Block or Warning**.
9. Select **OK** or **Apply**.

To create a FortiClient profile - CLI:

This example creates a profile for Windows and Mac computers.

```
config endpoint-control profile
  edit ep-profile1
    set device-groups mac windows-pc
    config forticlient-winmac-settings
      set forticlient-av enable
      set forticlient-wf enable
      set forticlient-wf-profile default
    end
  end
```

Support FortiClient for Linux

FortiClient for Linux (Ubuntu, CentOS, Red Hat, and Fedora) is also supported.

Syntax

```
config forticlient-winmac-settings
  config forticlient-operating-system
    edit <id>
      set os-type {ubuntu-linux | centos-linux | redhat-linux | fedora-linux | ...}
    next
    set forticlient-linux-ver <forticlient-version>
  end
```

Enforcing FortiClient registration

When you enable FortiTelemetry (formerly known as FortiHeartbeat) on an interface, the option to enforce FortiClient registration becomes available. Devices connecting to that interface are forced to register to the FortiGate and install FortiClient before gaining access to network services.

The following example includes editing the default FortiClient Profile to enforce real time antivirus protection and malicious website blocking.

To enforce FortiClient registration on the internal interface - GUI:

1. On the FortiGate, go to **System > Feature Visibility** and enable **Endpoint Control**.
2. Go to **Network > Interfaces** and edit the internal interface.
3. Under **Administrative Access**, enable **FortiClient Telemetry**.
4. Under **Admission Control**, enable **Enforce FortiClient Compliance Check**.
Once this is enabled, you have the option to **Exempt Sources** and/or **Exempt Destinations/Services**. If you were to exempt a source device, that device would not require FortiClient registration to access network services or the Internet.
5. Go to **Security Profiles > FortiClient Profiles**.
6. Under the **Security Posture Check**, enable **Realtime Protection, Up-to-date signatures**.

Endpoint compliance checking

Previously, as part of the Endpoint Compliance - Authorized Machine Detection feature, the administrator could specify a process name and SHA256 signature for a process, and only allow access to hosts with the specified

process/application running. The FortiGate verifies if the process name and hash is matched on the connecting host to allow access.

In FortiOS 6.0, however, the FortiGate only matches the process name, and matching the SHA256 signature is optional (since the process may be updated dynamically and the signature may not match). The administrator can specify a process name and not specify a checksum, and so only the file name will be matched. If both file name and MD5 are specified, then both fields will still be matched.

A host check table has been added to the FortiClient Profile GUI, which is similar to a policy table; the match is performed from top to bottom. At the bottom of the table, there is an implicit entry, representing everything that does not match the higher entries. This implicit entry is always available, but the administrator can change the action to either **present** or **absent** (in reference to the specified process/application).

Syntax

A new attribute `application-check-rule` determines if the entry is for checking the presence or absence of an application:

```
config endpoint-control profile
edit <name>
  config forticlient-winmac-settings
  ....
  config forticlient-running-app
  edit 1
    set app-name "MSOffice"
    set application-check-rule {present | absent}
    set process-name "word.exe"
  next
  ...
```

In addition, the `app-sha256-signature` entry is no longer mandatory, so long as the `process-name` entry is set:

```
config endpoint-control profile
edit <name>
  config forticlient-winmac-settings
  ....
  config forticlient-running-app
  edit <name>
    set app-name <name>
    set application-check-rule present
    set process-name "word.exe"
    set app-sha256-signature '' <== this field can be left empty
    set process-name2 "excel.exe"
    set app-sha256-signature2 '' <== not mandatory if process-name entry is set
    set process-name3 ''
    set app-sha256-signature3 '' <== not mandatory if process-name entry is set
    set process-name4 ''
    set app-sha256-signature4 '' <== not mandatory if process-name entry is set
  next
  ...
```

Enforcing FortiClient EMS requirements

FortiClient Compliance Profiles allow you to add up to three Enterprise Management Server (EMS) servers under **Security Profiles > FortiClient Compliance Profiles**.

This replaces the feature-related configuration (i.e AV, WF configuration) for compliance checks. Instead, if a FortiClient endpoint is managed by the defined EMS and is "in-sync" with the EMS profile then it is considered compliant.

An endpoint is considered compliant (thus allowed network access) only when the following conditions are met:

- the endpoint has FortiClient software
- the FortiClient software is managed by the authorized EMS server

Any endpoint that does not meet the above criteria (unless exempted) will be blocked from network access, regardless of FortiClient settings on that endpoint.

Syntax

```
config endpoint-control profile
  edit <name>
    config {forticlient-winmac-settings | forticlient-android-settings | forticlient-
      ios-settings}
      set forticlient-ems-compliance {enable | disable}
      set forticlient-ems-compliance-action {block | warning}
      set forticlient-ems-entries [addr1] [addr2] [addr3]
    next
  end
end

config endpoint-control settings
  set forticlient-ems-rest-api-call-timeout <milliseconds>
end
```

Changing the FortiClient installer download location

By default, FortiClient installers are downloaded from the FortiGuard network. You can also host these installers on a server for your users to download. In that case, you must configure FortiOS with this custom download location. For example, to set the download location to a customer web server with address custom.example.com, enter the following command:

```
config endpoint-control settings
  set download-location custom
  set download-custom-link "http://custom.example.com"
end
```

Storing FortiClient configuration files

Advanced FortiClient configuration files of up to 32k may be stored:

1. Enable the advanced FortiClient configuration option in the endpoint profile:

```
config endpoint-control profile
  edit "default"
    set forticlient-config-deployment enable
    set fct-advanced-cfg enable
    set fct-advanced-cfg-buffer "hello"
    set forticlient-license-timeout 1
    set netscan-discover-hosts enable
  next
end
```

2. Export the configuration from FortiClient (xml format).
3. Copy the contents of the configuration file and paste in the advanced FortiClient configuration box.

If the configure file is greater than 32k, you need to use the following CLI:

```
config endpoint-control profile
  edit <profile>
    config forticlient-winmac-settings
      config extra-buffer-entries
        edit <entry_id>
          set buffer xxxxxx
        next
      end
    end
  next
end
```

Blocking access to unsupported FortiClient endpoints

You can use the following command to deny registration of unsupported FortiClient endpoints. An unsupported FortiClient endpoint means the endpoint is running FortiClient but for some reason not all of the criteria are available to identify the endpoint, or the endpoint may be running an unsupported version of FortiClient. Information required that is not available could include the endpoint's IP address or MAC address is not visible.

```
config endpoint-control setting
  set forticlient-dereg-unsupported-client enable
end
```

Configuring the FortiClient offline grace period

Administrators can configure an offline grace period for registered and offline FortiClients so that PROBE can be processed and, as a result, endpoint compliance is not triggered.

- The grace period is allowed for a client that is compliant, registered, and offline.
- The grace period has a used status which determines if the client is before, during, or after grace period.
- Online and compliant clients will reset the grace status to unused.

Syntax

```
config endpoint-control settings
  set forticlient-offline-grace {enable | disable}
  set forticlient-offline-grace-interval <seconds>  <-- The default is 120
end
```

Configuring endpoint registration over a VPN

FortiGate units can register FortiClient-equipped endpoints over either an interface-based IPsec VPN or a tunnel-mode SSL VPN. After the user authenticates, the FortiGate unit sends the FortiClient application the IP address and port to be used for registration. If the user accepts the FortiGate invitation to register, registration proceeds and the FortiClient profile is downloaded to the client.

Users without FortiClient Endpoint Security connecting to the SSL VPN through a browser are redirected to a captive portal to download and install the FortiClient software.

Endpoint registration on an IPsec VPN

You can enable endpoint registration when you configure the FortiClient VPN or you can enable it on an existing FortiClient VPN.

To enable endpoint registration while configuring the VPN

- Enable **Allow Endpoint Registration** on the Policy & Routing page of the VPN Wizard when creating the FortiClient VPN.



This is only available when **Template Type** is set to **Remote Access** with a FortiClient **Remote Device Type**.

To enable endpoint registration on an existing VPN

1. Go to **Network > Interfaces** and edit the VPN's tunnel interface.
The tunnel is a virtual interface under the physical network interface.
2. In **Admission Control**, enable **FortiClient Telemetry**.
Optionally, you can also enable **Enforce FortiClient Telemetry for all FortiClients**. This forces endpoints to register with FortiClient before they have network access.
3. Select **OK**.

Endpoint registration on an SSL-VPN

To enable endpoint registration on the SSL-VPN

1. Go to **VPN > SSL-VPN Settings**.
2. In **Tunnel Mode Client Settings**, make sure **Allow Endpoint Registration** is enabled.
3. Select **Apply**.
4. Go to **Network > Interfaces** and edit the **ssl.root** interface.
5. In **Admission Control**, enable **FortiTelemetry**.
Optionally, you can also enable **Enforce FortiClient Telemetry for all FortiClients**. This forces endpoints to register with FortiClient before they have network access.
6. Select **OK**.

This procedure does not include all settings needed to configure a working SSL-VPN.

Synchronizing endpoint registrations

To support roaming users in a network with multiple FortiGate units, you need to configure synchronization of the endpoint registration databases between the units. The registered endpoints are then recognized on all of the FortiGate units. This is configured in the CLI. For example, to synchronize this FortiGate unit's registered endpoint database with another unit named `other1` at IP address 172.20.120.4, enter:

```
config endpoint-control forticlient-registration-sync
  edit other1
    set peer-ip 172.20.120.4
  end
```

Assigning FortiClient Profiles using Microsoft AD user groups

When FortiClient Telemetry connects to FortiGate, the user's AD domain name and group are sent to FortiGate. Administrators may configure FortiGate to assign Endpoint Profiles based on the end user's AD domain group membership.

The following steps are discussed in more detail:

- [Configuring users and groups on AD servers](#)
- [Configuring FortiAuthenticator](#)
- [Configuring FortiGate](#)
- [Connecting FortiClient Telemetry to FortiGate](#)
- [Monitoring FortiClient connections](#)

Configuring users and groups on AD servers

Create the user accounts and groups on the AD server. Groups may have any number of users. A user may belong to more than one group at the same time.

Configuring FortiAuthenticator

Configure FortiAuthenticator to use the AD server you created. See the FortiAuthenticator Administration Guide in the [Fortinet Document Library](#).

Configuring FortiGate

FortiGate

Add the FortiAuthenticator or Fortinet Single Sign-On Agent (FSSO):

1. Go to **Security Fabric > Fabric Connectors**.
2. Select **Create New** in the toolbar. The New Fabric Connector window opens.
3. Under **SSO/Identity**, select **Fortinet Single-Sign-On Agent**.
4. Enter the information required for the agent. This includes the name, primary and (optional) secondary IP addresses, and passwords. Select More FSSO agents to add up to three additional agents.
5. For **Collector Agent AD access mode**, select **Standard** or **Advanced**.
 - a. **Standard**: select Users/Groups to include as Single-Sign-On accounts.
 - b. **Advanced**: select an LDAP server in the dropdown list.
6. Select **OK** to save the agent configuration.

Create a user group:

1. Go to **User & Device > User Groups**.
2. Select **Create New** in the toolbar. The New User Group window opens.
3. In the **Type** field, select **Fortinet Single-Sign-On (FSSO)**.
4. Select members from the dropdown list.
5. Select **OK** to save the group configuration.

Configure the FortiClient profile:

1. Go to **Security Profiles > FortiClient Compliance**.
2. Select Create New in the toolbar. The New FortiClient Profile window opens.
3. Enter a profile name and optional comments.
4. In the **Assign Profile To** dropdown list select the FSSO user group(s).
5. Configure FortiClient configuration as required.
6. Select **OK** to save the new FortiClient profile.



Create any number of FortiClient profiles with different groups and different settings. The default profile will be assigned to users who connect successfully, but have no matching FortiClient profile.

Configure the firewall policy:

Configure the firewall policy. Ensure Compliant with FortiClient Profile is selected in the policy.

Connecting FortiClient Telemetry to FortiGate

The Microsoft Windows system where FortiClient is installed should join the domain of the AD server configured earlier. Users may log in with their domain username.

Following this, endpoint connections send the logged-in user's name and domain to the FortiGate. The FortiGate will assign the appropriate profiles based on the configurations.

Monitoring FortiClient connections

The following FortiOS CLI command lists information about connected clients. This includes domain-related details for the client if any.

```
diagnose endpoint record-list
Record #1:
  IP_Address = 172.172.172.111 (1)
  MAC_Address = b0:ac:6f:70:e0:a0
  Host MAC_Address = b0:ac:6f:70:e0:a0
  MAC list = b0-ac-6f-70-e0-a0;
  VDOM = root
  Registration status: Forticlient installed but not registered
  Online status: offline
  DHCP on-net status: off-net
  DHCP server: None
  FCC connection handle: 6
  FortiClient version: 5.1.29
  AVDB version: 22.137
  FortiClient app signature version: 3.0
  FortiClient vulnerability scan engine version: 1.258
  FortiClient feature version status: 0
  FortiClient UID: BE6B76C509DB4CF3A8CB942AED2064A0 (0)
  FortiClient config dirty: 1:1:1
  FortiClient KA interval dirty: 0
  FortiClient Full KA interval dirty: 0
  FortiClient server config: d9f86534f03fbed109676ee49f6cfc09::
  FortiClient config: 1
  FortiClient iOS server mconf:
```

```

FortiClient iOS mconf:
FortiClient iOS server ipsec_vpn mconf:
FortiClient iOS ipsec_vpn mconf:
Endpoint Profile: Documentation
Reg record pos: 0
Auth_AD_groups:
Auth_group:
Auth_user:
Host_Name:
OS_Version: Microsoft Windows 7 , 64-bit Service Pack 1 (build 7601)
Host_Description: AT/AT COMPATIBLE
Domain:
Last_Login_User: FortiClient_User_Name
Host_Model: Studio 1558
Host_Manufacturer: Dell Inc.
CPU_Model: Intel(R) Core(TM) i7 CPU Q 720 @ 1.60GHz
Memory_Size: 6144
Installed features: 55
Enabled features: 21
online records: 0; offline records: 1
status -- none: 0; uninstalled: 0; unregistered: 1; registered: 0; blocked: 0

```

Modifying the endpoint protection replacement messages

If the security policy has **Redirect all non-compliant/unregistered FortiClient compatible devices to a captive portal** enabled, users of non-compliant devices are redirected to a captive portal that is defined by the **Endpoint NAC Download Portal** replacement message. There are different portals for Android, iOS, Mac, Windows, Quarantine, and “other” devices.

To modify the the endpoint protection replacement messages

1. Go to **System > Replacement Messages** and select **Extended View**.
2. In the **Endpoint Control** section select the message that you want to edit.
The replacement message and its HTML code appear in a split screen in the lower half of the page.
3. Modify the text as needed and select **Save**.

Monitoring endpoints

Go to **Monitor > FortiClient Monitor** to monitor endpoints.

The **Monitor** page allows the user to view FortiClient endpoint devices grouped by interface and then sub-grouped by compliance status. Compliance status can be compliant, non-compliant, exempt, or quarantined.

Status	Enforcement Enabled	Enforcement Disabled
Compliant	List only active FortiClient endpoints.	No devices listed.
Not-compliant	List devices not-compliant with FortiClient profile, so long as they are not exempt.	No devices listed.

Status	Enforcement Enabled	Enforcement Disabled
Exempt*	List FortiClient endpoints exempt from FortiClient compliance.	List of all user devices except those quarantined by the administrator.
Quarantined	List devices quarantined by the administrator.	List devices quarantined by the administrator.

* Includes device exempt reasons as any combination of device, device category/group, and source address.

You can see the reasons for non-compliance by right-clicking on an endpoint in the list.

Proxy options

Certain inspections defined in security profiles require that the traffic be held in proxy while the inspection is carried out. When a security profile requiring the use of a proxy is enabled in a policy, the **Proxy Options** field is displayed. The Proxy Options define the parameters of how the traffic will be processed and to what level the traffic will be processed. There can be multiple security profiles of a single type. There can also be a number of unique Proxy Option profiles. As the requirements for a policy differ from one policy to the next, a different Proxy Option profile for each individual policy can be configured or one profile can be repeatedly applied.

The **Proxy Options** refer to the handling of the following protocols:

- HTTP
- SMTP
- POP3
- IMAP
- FTP
- NNTP
- MAPI
- DNS

The configuration for each of these protocols is handled separately.

The use of different proxy profiles and profile options

Just like other components of the FortiGate, different Proxy Option profiles can be configured to allow for granular control of the FortiGate. In the case of the Proxy Option profiles the thing that you will want to focus on is the matching up of the correct profile to a firewall policy that is using the appropriate protocols. If you are creating a Proxy Option profile that is designed for policies that control SMTP traffic into your network you only want to configure the settings that apply to SMTP. You do not need or want to configure the HTTP components.

Proxy Options profile components

Highlighted below are certain features available in the **Proxy Options** security profile.

Log Oversized Files

This setting enables logging of the occurrence of oversized files being processed. It does not change how they are processed. It only enables the FortiGate unit to log that they were either blocked or allowed through. A common practice is to allow larger files through without antivirus processing. This allows you to get an idea of how often this happens and decide on whether or not to alter the settings relating to the treatment of oversized files.

The setting of the threshold for oversized files and emails is found on the Security Profiles > Proxy Options page under **Common Options**.

RPC over HTTP

FortiGate units with firmware version 5.4 and higher support RPC over HTTP. This protocol is used by the Microsoft Exchange Server to perform virus scanning of Microsoft Exchange Server email that uses RPC over HTTP. To enable this feature, go to **Security Profiles > Proxy Options** and enable **RPC over HTTP**.

Protocol Port Mapping

To optimize the resources of the unit, the mapping and inspection of protocols can be enabled or disabled.

Each of the protocols listed in the GUI has a commonly used default TCP port, however, the port used by the protocols can be individually modified. It can also be set to inspect any port with flowing traffic for that particular protocol. The headers of the packets indicate which protocol generated the packet.

Comfort Clients

When proxy-based antivirus scanning is enabled, the FortiGate unit buffers files as they are downloaded. Once the entire file is captured, the FortiGate unit begins scanning the file. During the buffering and scanning procedure, the user must wait. After the scan is completed, if no infection is found, the file is sent to the next step in the process flow. If the file is a large one this part of the process can take some time. In some cases enough time that some users may get impatient and cancel the download.

The **Comfort Clients** feature mitigates this potential issue by feeding a trickle of data while waiting for the scan to complete. The user then knows that processing is taking place and that there hasn't been a failure in the transmission. The slow transfer rate continues until the antivirus scan is complete. Once the file has been successfully scanned and found to be clean of any viruses, the transfer will proceed at full speed.

If there is evidence of an infection, the FortiGate unit caches the URL and drops the connection. The client does not receive any notification of what happened because the download to the client had already started. Instead, the download stops and the user is left with a partially downloaded file. If the user tries to download the same file again within a short period of time, the cached URL is matched and the download is blocked. A notification that the download has been blocked is displayed. The number of URLs in the cache is limited by the size of the cache.

Client comforting is available for HTTP and FTP traffic. If your FortiGate unit supports SSL content scanning and inspection, you can also configure client comforting for HTTPS and FTPS traffic.



Buffering the entire file allows the FortiGate unit to eliminate the danger of missing an infection due to fragmentation because the file is reassembled before examination. Client comforting can send unscanned and therefore potentially infected content to the client. You should only enable client comforting if you are prepared to accept this risk. Keeping the client comforting interval high and the amount low will reduce the amount of potentially infected data that is downloaded.

Block Oversized File/Email

This feature is related to antivirus scanning. The FortiGate unit has a finite amount of resources that can be used to buffer and scan a file. If a large file such as an ISO image or video file was to be downloaded this could overwhelm or exceed the memory of the FortiGate, especially if there were other large files being downloaded at the same time. For this reason, the treatment of large files needs to be addressed.

A threshold is assigned to identify an oversized file or email. This can be set at any size from 1 MB to 10 MB. Any file or email over this threshold will not be processed by policies applying the Antivirus security profile.



It should be noted that in terms of probability that malware is more likely to be found in smaller files than in larger files. A number of administrators take this into account when they lower the default threshold so as to lessen the impact on memory if they see the FortiGate unit going into conserve mode on a regular basis.

Chunked Bypass

The HTTP section allows the enabling of **Chunked Bypass**. This refers to the mechanism in version 1.1 of HTTP that allows a web server to start sending chunks of dynamically generated output in response to a request before actually knowing the actual size of the content. Where dynamically generated content is concerned, enabling this feature means that there is a faster initial response to HTTP requests. From a security stand point, enabling this feature means that the content will not be held in the proxy as an entire file before proceeding.

Allow Fragmented Messages

The specifications of RFC 2046 allow for the breaking up of emails and sending the fragments in parallel to be rebuilt and read at the other end by the mail server. It was originally designed to increase the performance over slower connections where larger email messages were involved. It will depend on your mail configuration if this is even possible for your network but outside of Microsoft Outlook and Outlook Express, not many email clients are set up to break up messages like this. The drawback of this feature is that if malware is broken up between multiple fragments of the message the risk is run that it will not be detected by some antivirus configurations because the code may not all be present at the same time to identify.

Append Email Signature

The **Append Email Signature** feature ensures that all of the emails going out of a particular network has the appropriate signature or corporate message, for example. These appended emails do not replace existing signatures.

Examples could include things like:

- Without prior approval the email should not be forwarded.
- Please be environmentally friendly and don't print out emails
- For questions regarding the purchasing of our products please call...

It can be anything that the organization would like as long as it is in text format. The use of this feature usually works best in an environment where there is some standardization of what goes into the personal signatures of the senders so that there is no duplication or contradiction of information in the signatures.

SSL/SSH inspection

Individual deep inspection security profiles can be created depending on the requirements of the policy. Depending on the inspection profile selected, you can:

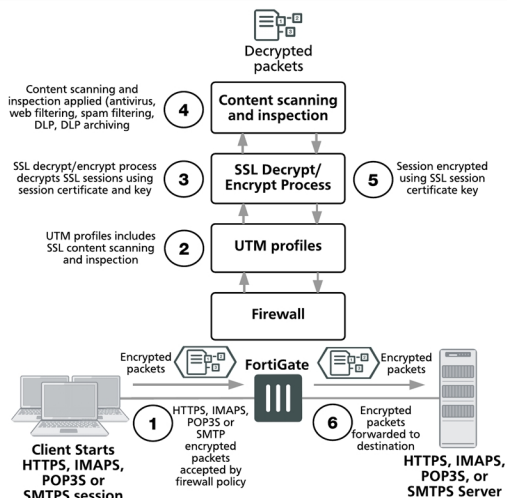
- Configure which Certificate Authority (CA) certificate will be used to decrypt the Secure Sockets Layer (SSL) encrypted traffic.
- Configure whether a specific SSL protocol will be inspected, blocked or bypassed.
- Configure which ports will be associated with which SSL protocols for the purpose of inspection.
- Configure which websites or website categories will be exempt from SSL inspection
- Identify how to treat invalid, unsupported or untrusted SSL certificates.
- Determine which inspection method will be applied to Secure Shell (SSH) / SSL traffic.

SSL inspection

Secure Sockets Layer (SSL) content scanning and inspection allows you to apply antivirus scanning, web filtering, FortiGuard Web Filtering, and email filtering to encrypted traffic. To perform SSL content scanning and inspection, the FortiGate unit does the following:

- intercepts and decrypts HTTPS, IMAPS, POP3S, SMTPS, and FTPS sessions between clients and servers (FortiGate SSL acceleration speeds up decryption)
- applies content inspection to decrypted content, including:
 - HTTPS, IMAPS, POP3S, and SMTPS Antivirus, DLP, and DLP archiving
 - HTTPS web filtering and FortiGuard web filtering
 - IMAPS, POP3S, and SMTPS email filtering
- encrypts the sessions and forwards them to their destinations.

FortiGate SSL content scanning and inspection packet flow



SSL inspection and privacy

Normally, SSL decrypted content is temporarily stored in system memory for content scanning. If Malware is found the infected content is deleted and a message is sent to the destination instead. If no Malware is found the content is re-encrypted and forwarded to its destination. Administrators are not able to access or view the decrypted content.

There are two exceptions that you should be aware of if you have privacy concerns:

- If Sandbox inspection is enabled, either with an on-premises FortiSandbox device or FortiCloud Sandbox, decrypted files can be sent to FortiSandbox or FortiSandbox cloud where they can be viewed by system administrators.
- For flow-based SSL inspection, if SSL mirroring is enabled it is possible to "mirror" or send a copy of the decrypted content to one or more FortiGate interfaces so that the content can be collected by a raw packet capture tool for archiving or analysis. This feature is only available if the inspection mode is set to flow-based.



Decryption, storage, inspection, and use decrypted content is subject to local privacy rules. Use of these features could enable malicious users with administrative access to your FortiGate to harvest sensitive information submitted using an encrypted channel.

For increased privacy of sensitive information, you can use the SSL inspection exemptions feature, described below, to exempt sensitive communication from decryption.

SSL inspection exemptions

When you are using a browser to visit SSL encrypted sites and are using a certificate that does not match the certificate of the site, you are presented with a warning message and the option of continuing with the untrusted certificate, or terminating the session. However, there are a number of applications that use SSL encrypted traffic. Some applications will not allow SSL traffic that isn't signed with a trusted certificate. These applications do not necessarily give the option to manually indicate that we trust the certificate or the site. If the option is available, the customer may choose to import needed SSL certificates into Local Certificates and configure a policy for communication for that application.

To assist in preventing loss of access to these sites while still enabling the SSL inspection of the rest of the internet traffic, a method of exempting either web categories or specific sites has been developed. To exempt a large group of sites, the **SSL/SSH Inspection** profile can be configured to exempt FortiGuard Categories. There are three preselected categories due to the high likelihood of issues with associated applications with the type of websites included in these categories.

- Finance and Banking
- Health and Wellness
- Personal Privacy

Other more specific websites can be added to the exemption list by going to **Security Profiles > SSL Inspection**, selecting the appropriate profile, and adding addresses under **Exempt from SSL Inspection**.



When you create a custom web category and tell the inspection profile to exempt that category, you may find some URLs in that category are still inspected. As a best practice, use the [Static URL filter](#) "Exempt" option instead.

Your FortiGate unit has two pre-configured SSL/SSH Inspection profiles that cannot be edited: **certificate-inspection** and **deep-inspection**. You must clone and edit the pre-configured profiles or create a new profile to exempt any additional sites or FortiGuard categories.

Allow Invalid SSL Certificates

It might seem like a straightforward decision that the allowing of invalid SSL certificates must be bad and, therefore, should not be allowed. However, there can be some reasons that applying this feature should be considered.

At a purely technical level, a properly formed certificate will encrypt the data so that it can only be read by the intended parties and not be read by anyone sniffing traffic on the network. For this reason, people will often use self-signed certificates. These self-signed certificates are free and will encrypt the data just as securely as a purchased certificate. The self-signed certificates, however, are not likely to be recognized by the CA certificate store so will be considered by any checks against that store as invalid.

On the other hand, one of the services the vendors provide is verification of identity of those that purchase their certificates. This means that if you see a valid certificate from a site that identified itself as being from “valid-company.com” that you can be reasonably sure that the site does belong to that company and not a false site masquerading as being part of that company.

You can allow invalid SSL certificates by going to **Security Profiles > SSL Inspection**, selecting the appropriate profile, and enabling **Allow Invalid SSL Certificates**.

During the SSL handshake, a number of checks are made to verify the validity of the certificate.

One source of the checks, is against a CA certificate store inside FortiOS. This is the same CA bundle used by the browser Mozilla Firefox.

Updates to the store are:

- With each new version of FortiOS
- Via internal FGD
- Possible with some builds via FTP

Details of the CA certificate store can be found at: <https://curl.haxx.se/docs/caextract.html>

The following checks are made for validity:

Validity Check	Description
Signature	One of the things being checked against the CA bundle is the certificate signature. These signatures are generated via directly signing by the CA's private key.
Expiration date	All certificates have an expiry date. The date, based on the device's clock/calendar is compared to the expiry date of the certificate.
Revoked list	Periodically, certificates are revoked. If a certificate has been revoked it is put on a list. Whenever a certificate is being verified, it is checked against this list.

Validity Check	Description
Self signed certificate	In the case of self-signed certificates, the IPS engine and proxy have different handling. IPS engine will keep and use the certificate self-signed certificate, but the public key will be replaced so that SSL inspection can take place. The proxy engine will re-sign the certificate with the untrusted CA certificate. The mechanics are similar but the net effect for the user is similar. The user will get warnings from browsers. The users can choose to remember the self-signed certificate in some browsers, but cannot do the same thing with the certificate re-signed with the untrusted CA.
Intermediate CA with a weak hash algorithm, such as MD5, SHA1	<p>Some browsers like Chrome or Firefox will give a warning because of a weak signature algorithm (visit https://sha1-intermediate.badssl.com to test).</p> <p>In the IPS Engine, in order to convey the weak intermediate CA back to client, the signature hash algorithm is downgraded in the re-signed server certificate to the weakest algorithm used in the original certificate chain.</p> <p>In the Proxy Engine - In the case of a weak signature algorithm, the Proxy engine will treat the connection as untrusted, and re-sign the server certificate with the untrusted CA. The final user experience is different. Instead of a warning like "NET::ERR_CERT_WEAK_SIGNATURE_ALGORITHM" that you would get in Chrome, you will get a warning that the certificate couldn't be verified (because of the signing CA is not trusted or imported into the user's web browser).</p>

Flow-based behaviour

In flow-based mode, a certificate will be considered as **invalid** if it has expired.

In addition, a certificate will be considered as **untrusted** if one or more of the following conditions are met:

- If the chain is broken or incomplete.
- If it is part of the CRL.
- If the CA certificate was not imported to the FortiGate, or it is not in the FortiGate CA certificate store.

Why use SSL inspection

Most of us are familiar with Hypertext Transfer Protocol Secure (HTTPS) and how it protects a variety of activities on the Internet by applying Secure Sockets Layer (SSL) encryption to the web traffic. However, there are risks associated with its use, since encrypted traffic can be used to get around your network's normal defenses.

For example, you might download a file containing a virus during an e-commerce session. Or you could receive a phishing email containing a seemingly harmless downloader file that, when launched, creates an encrypted session to a command and control (C&C) server and downloads malware onto your computer. Because the sessions in these attacks are encrypted, they might get past your network's security measures.

To protect your network from these threats, SSL inspection is the key your FortiGate uses to unlock encrypted sessions, see into encrypted packets, find threats, and block them. SSL inspection not only protects you from attacks that use HTTPS, but also from other commonly used SSL-encrypted protocols, such as SMTPS, POP3S, IMAPS, and FTPS.

Full SSL inspection

To make sure that all SSL encrypted content is inspected, you must use full SSL inspection (also known as deep inspection). When full SSL inspection is used, the FortiGate impersonates the recipient of the originating SSL session, then decrypts and inspects the content. The FortiGate then re-encrypts the content, creates a new SSL session between the FortiGate and the recipient by impersonating the sender, and sends the content to the sender.

When the FortiGate re-encrypts the content it uses a certificate stored on the FortiGate. The client must trust this certificate to avoid certificate errors. Whether or not this trust exists depends on the client, which can be the computer's OS, a browser, or some other application, which will likely maintain its own certificate repository. For more information about this, see the recipe [Preventing certificate warnings](#) on the [Fortinet Cookbook](#) site.

There are two deployment methods for full SSL inspection:

1. Multiple Clients Connecting to Multiple Servers:

- Uses a CA certificate (which can be uploaded using the **Certificates** menu).
- Typically applied to outbound policies where destinations are unknown (i.e. normal web traffic).
- Address and web category whitelists can be configured to bypass SSL inspection.

2. Protecting SSL Server

- Uses a server certificate (which can be uploaded using the **Certificates** menu) to protect a single server.
- Typically used on inbound policies to protect servers available externally through Virtual IPs
- Since this is typically deployed "outside-in" (clients on the Internet accessing server(s) on the internal side of the FortiGate), server certificates using the public FQDN of the server are often purchased from a commercial Certificate Authority and uploaded to the FortiGate. This avoids client applications generating SSL certificate errors due to certificate mismatch.

More detail is available in the Fortinet Knowledge Base. Check these technical notes:

- [How to Enable SSL inspection from the CLI and Apply it to a Policy](#)
- [How to block web-based chat on Gmail webmail using App Sensor + SSL inspection](#)

SSL certificate inspection

FortiGates also support a second type of SSL inspection, called SSL certificate inspection. When certificate inspection is used, the FortiGate only inspects the header information of the packets.

Certificate inspection is used to verify the identity of web servers and can be used to make sure that HTTPS protocol isn't used as a workaround to access sites you have blocked using web filtering.

The only security feature that can be applied using SSL certificate inspection mode is web filtering. However, since only the packet is inspected, this method does not introduce certificate errors and can be a useful alternative to full SSL inspection when web filtering is used.

Troubleshooting

The most common problem with SSL inspection is users receiving SSL errors when the CA certificate is not trusted. This is because by default the FortiGate uses a certificate that is not trusted by the client. There are two ways to fix this:

1. All users must import the FortiGate's default certificate into their client applications as a trusted certificate.
2. Configure the FortiGate to use a certificate that is already trusted by your clients. For example, a certification signed by a CA that your clients already trust.

The first method can be more labor intensive because you have to distribute a certification to all clients. This can also be an ongoing problem as new clients are added to your network. The second method is usually less work but may require paying for a CA. Both of these methods are covered in the recipe [Preventing Certificate Warnings](#).

If you choose to install the certificate on client applications, this can be done with greater ease in a Microsoft Active Directory domain environment by using Group Policy Objects to install the certificate on domain members. Check that the Group Policy has propagated to all computers by opening Internet Explorer on a workstation PC, opening **Tools > Internet Options > Content > Certificates > Trusted Root Certification Authorities**, and ensuring that the FortiGate's certificate is present.

For corporate-owned mobile devices, MDM solutions like AirWatch, MobileIron, or Fiberlink, use Simple Certificate Enrollment Protocol (SCEP) to ease certificate enrollment.

Best practices

Because all traffic needs to be decrypted, inspected, and re-encrypted, using SSL inspection can reduce overall performance of your FortiGate. To make sure you aren't using too many resources for SSL inspection, do the following:

- **Know your traffic** – Know how much traffic is expected and what percent of the traffic is encrypted. You can also limit the number of policies that allow encrypted traffic.
- **Be selective** – Use white lists or trim your policy to apply SSL inspection only where it is needed.
- **Use hardware acceleration** - FortiGate models with either the CP6 or CPU processor have an SSL/TLS protocol processor for SSL content scanning and SSL acceleration. For more information about this, see the [Hardware Acceleration handbook](#).
- **Test real-world SSL inspection performance yourself** - Use the flexibility of FortiGate's security policy to gradually deploy SSL inspection, rather than enabling it all at once.

Creating or editing an SSL/SSH Inspection profile

1. Go to **Security Profiles > SSL/SSH Inspection**. This will open to one of the existing profiles. Your FortiGate unit has two pre-configured SSL/SSH Inspection profiles that cannot be edited: certificate-inspection and deep-inspection. You must clone and edit the pre-configured profiles or create a new profile to exempt any additional sites or FortiGuard categories. The links for the actions are located in the upper right hand corner of the window.
 - To view a list of the existing profiles select the List icon (a page) at the far right.
 - To clone an existing profile, select the Clone icon (one page behind another), second from the right
 - To create a new profile, select the Create New icon ("+" symbol), third from the right.
 - To view or edit an existing profile, choose it from the dropdown menu field.
2. **Name Field:**
Give the profile an easily identifiable name that references its intent.
3. **Comments Field:**
Enter any additional information that might be needed by administrators, as a reminder of the profile's purpose and scope.
4. **SSL Inspection Options:**
 - a. Enable SSL Inspection of:

- Multiple Clients Connecting to Multiple Servers - Use this option for generic policies where the destination is unknown.
- Protecting SSL Server - Use this option when setting up a profile customized for a specific SSL server with a specific certificate.
- b. Inspection Method
The options here are:
 - SSL Certificate Inspection - only inspects the certificate, not the contents of the traffic.
 - Full SSL Inspection - inspects all of the traffic.
- c. CA Certificate
Use the drop down menu to choose which one of the installed certificates to use for the inspection of the packets or click on **Download Certificate**.
- d. Untrusted SSL Certificates
Select an action for untrusted SSL certificates.
- d. **Protocol Port Mapping** / Inspect All Ports
Enable the ability to inspect all ports by checking the box. If the feature is not enabled, specify in the field next to the listed protocols, the port through which that protocols traffic will be inspected. Traffic of that protocol going through any other port will not be inspected.



If you select **Inspect All Ports**, then only the IPS engine is used for inspection.

5. Exempt from SSL Inspection:

Use the dropdown menus in this section to specify any reputable websites, FortiGuard Web Categories, or addresses will be exempt from SSL inspection.

- Reputable Websites - Enable this option to exempt any websites identified by FortiGuard as reputable.
- Web Categories - By default the categories of Finance and Banking, Health and Wellness, and Personal Privacy, have been added as these are one that are most likely to have applications that will require a specific certificate.
- Addresses - These can be any of the Address objects that have an interface of "Any".
- Log SSL exemptions - Enable this option to log all SSL exemptions

6. SSH Inspection Options:

a. SSH Deep Scan

Toggle to disable or enable the feature

b. SSH Port

The available options are:

- **Any** - choosing this option will search all of the traffic regardless of service or TCP/IP port for packets that conform to the SSH protocol
- **Specify** - choosing this option will restrict the search for SSH protocol packets to the TCP/IP port number specified in the field. This is not as comprehensive but it is easier on the performance of the firewall.
- d. Protocol Actions
 - Exec - Block, Log or neither. Select using check boxes.
 - Port-Forward - Block, Log or neither. Select using check boxes.

- SSH-Shell - Block, Log or neither. Select using check boxes.
- X11-Filter - Block, Log or neither. Select using check boxes.

6. Common Options:

- Allow Invalid SSL Certificates
Check the box to enable the passing of traffic with invalid certificate
- Log SSL anomalies
Check the box to allow the Logging function to record traffic sessions containing invalid certificates



The **Full SSL Inspection** method is enabled by default when creating a new SSL/SSH Inspection profile. There are situations where this feature can cause issues so be sure that you would like it enabled before applying the inspection profile.

Secure white list database

You can enable a feature that gathers a list of reputable domain names that can be excluded from SSL deep inspection. This list is periodically updated and downloaded to FortiGate units through FortiGuard.

Go to **Security Profiles > SSL Inspection**, enable **Exempt from SSL Inspection**, and enable **Reputable Websites**. The reputable websites are rated by FortiGuard. Web Filtering.

CLI syntax:

```
config firewall ssl-ssh-profile
  edit deep-inspection
    set whitelist enable
  end
end
```

SSH MITM deep inspection

As vulnerabilities of OpenSSH continue to be exposed, it has become necessary to detect such attacks, which requires the ability to decrypt the SSH tunnel to check the data. This feature introduces comprehensive security controls on SSH Man-in-the-Middle (MITM) deep inspections, including:

- SSH filter profiles to control SSH tunnel types and filtering on SSH shell commands.
- SSH proxy policies to apply a proxy firewall policy with user authentication on SSH session.
- Support for SSH tunnel policy to perform access control for TCP/IP port forwarding traffic that is tunneled through the SSH proxy. IPS scanning can be applied to the tunneled traffic.
- Support for SSH trust to detect and prevent SSH MITM attacks.

Syntax

1. Add SSH related option in ssl-ssh-profile for proxy mode profile

- a. Add option to bypass or block unsupported SSH protocol (Deep scan only supports SSH 2.0)

```
config firewall ssl-ssh-profile
  edit <name>
    config ssh
      set unsupported-version {bypass | block}
    next
  end
```

```
end
```

b. Add option to enable SSH proxy policy check

```
config firewall ssl-ssh-profile
  edit <name>
    config ssh
      set ssh-policy-check {enable | disable}
      set ssh-tun-policy-check {enable | disable}
    next
  end
end
```



When SSH proxy policy check is enabled, proxy will check "SSH proxy" policy for SSH traffic and check "SSH tunnel" policy for TCP/IP port forwarding traffic.

c. Move block/log options for x11-filter/ssh-shell/exec/port-forward to SSH filter profile

2. SSH filter profile

a. Support options to block or log x11-filter/ssh-shell/exec/port-forward/sftp

```
config ssh-filter profile
  edit <name>
    set block {x11-filter | ssh-shell | exec | port-forward | sftp}
    set log {x11-filter | ssh-shell | exec | port-forward | sftp}
  next
end
```

b. Add Shell command filters

```
config ssh-filter profile
  edit <name>
    config shell-commands
      edit <id>
        set type {simple | regex}
        set pattern <cmd-string>
        set action {block | allow}
        set log {enable | disable}
        set alert {enable | disable}
        set severity {low | medium | high | critical}
      next
    end
    set default-command-log {enable | disable}
  end
```

3. Allow SSH filter profile to be set for config firewall policy when UTM is enabled.

4. Support SSH proxy policy for SSH sessions

a. Add a proxy type ssh into config firewall proxy-policy

```
config firewall proxy-policy
  edit <pol-id>
    set proxy ssh
  next
end
```

b. When user/user-group is set in SSH proxy policy, firewall authentication can be done for SSH proxy traffic. Authentication rule for SSH is added:

```

config authentication rule
  edit <name>
    set protocol ssh
  next
end

```

i. "Basic" authentication scheme:

```

config authentication scheme
  edit "ssh-active"
    set method basic
    set user-database "local" #or LDAP server
  next

```

ii. "ssh-publickey" authentication scheme:

```

config authentication scheme
  edit "ssh-pkey"
    set method ssh-publickey
    set user-database "local" #or LDAP server
    set ssh-ca "server-ca"
  next

```



User name is embedded in ssh-publickey. User group information will be retrieved if the publickey is validated by CA.

iii. Both "Basic" and "ssh-publickey" authentication scheme:

```

config authentication scheme
  edit "ssh-pkey"
    set method basic ssh-publickey
    set user-database "local" #or LDAP server
    set ssh-ca "server-ca"
  next

```

5. Support SSH tunnel policy to do access control for TCP/IP port forwarding traffic.

a. Add a proxy type ssh-tunnel into config firewall proxy-policy

```

config firewall proxy-policy
  edit <pol-id>
    set proxy ssh-tunnel
    set action {accept | deny}
  next
end

```

b. Support allow or deny and IPS sensor/app-control the traffic.

6. Support SSH trust to detect and prevent from SSH MITM attacks

a. Define trusted SSH hostkey for specific SSH server

```

config firewall ssh host-key
  edit <name>
    set status {trusted | revoked}
    set type {RSA | DSS | ECDSA}
    set nid <NID of ECDSA key>
    set ip <ip>
    set port <port>
    set hostname <name>
    set public-key <hostkey>
  next
end

```

b. Define trusted/untrusted CAs for hostkey signing. Any hostkey signed by trust CA is trusted unless the hostkey is revoked.

```
config firewall ssh local-ca
  edit <name>
    set password <passwd>
    set public-key <public key>
    set private-key <private key>
    set source {build-in | user}
  next
end
```



The system creates two build-in SSH CAs: Fortinet_SSH_CA and Fortinet_SSH_CA_Untrusted. The CAs are used to re-sign a server host key with local host-key using trusted/untrusted CA when the server host key is trusted or untrusted.

c. Define local hostkey templates for trusted re-signing. Be default, they are generated automatically.

```
config firewall ssh local-key
  edit <name>
    set password <passwd>
    set public-key <public key>
    set private-key <private key>
    set source {build-in | user}
  next
end
```



- i. The system creates different types of local host keys as default re-signing templates: Fortinet_SSH_RSA2048, Fortinet_SSH_DSA1024, Fortinet_SSH_ECDSA256, Fortinet_SSH_ECDSA384, Fortinet_SSH_ECDSA512, Fortinet_SSH_ED25519, Fortinet_SSH_RSA1024.
- ii. Admin can load their own local host keys and use them for MITM re-signing in config firewall ssh setting.

d. Per-VDOM SSH settings

```
config firewall ssh setting
  set caname <trusted-ca>
  set untrusted-caname <untrusted-ca>
  set hostkey-rsa <hostkey-rsa>
  set hostkey-dss <hostkey-dss>
  set hostkey-ecdsa256 <hostkey-ecdsa256>
  set hostkey-ecdsa384 <hostkey-ecdsa384>
  set ed25519-key <ed25519-key>
  set host-trusted-check {enable | disable}
end
```



- i. When a hostkey is trusted and signed by a CA, SSH proxy re-signs appropriate type of hostkey using trusted CA.
- ii. When a host is trusted but not signed, SSH proxy sends back appropriate type of hostkey.
- iii. When a hostkey is untrusted and signed by a CA, SSH proxy re-signs a temporary hostkey (1 hour life time) using untrusted CA.
- iv. When a host is trusted but not signed, SSH proxy sends back a temporary hostkey (one hour life time).

SSL server table for SSL offloading

An SSL server table can now be used for SSL offloading. This feature was introduced with the release of FortiOS 5.4.0.

CLI Syntax

```
config firewall ssl-ssh-profile
  edit <name>
    set use-ssl-server {enable|disable}
  next
end
```

Custom Application & IPS Signatures

Creating a custom IPS signature

The FortiGate predefined signatures cover common attacks. If you use an unusual or specialized application or an uncommon platform, add custom signatures based on the security alerts released by the application and platform vendors.

You can add or edit custom signatures using the GUI or the CLI.

To create a custom signature

1. Go to **Security Profiles > Intrusion Prevention**.
2. Select **[View IPS Signatures]**
3. Select **Create New** to add a new custom signature.
4. Enter a **Name** for the custom signature.
5. Enter the **Signature**. For information about completing this field, see [Custom signature syntax](#) and [Custom signature keywords](#).
6. Select **OK**.

Custom signature syntax

All custom signatures follow a particular syntax. Each begins with a header and is followed by one or more keywords. A custom signature definition is limited to a maximum length of 512 characters. A definition can be a single line or span multiple lines connected by a backslash (\) at the end of each line.

A custom signature definition begins with a header, followed by a set of keyword/value pairs enclosed by parenthesis [()]. The keyword and value pairs are separated by a semi colon (;) and consist of a keyword and a value separated by a space. The basic format of a definition is HEADER (KEYWORD VALUE;)

You can use as many keyword/value pairs as required within the 512 character limit. To configure a custom signature, go to **Security Profiles > Intrusion Prevention**, select **View IPS Signatures**, select **Create New**, and enter the data directly into the **Signature** field, following the guidance in the next topics.

The table below shows the valid characters and basic structure. For details about each keyword and its associated values, see [Custom signature keywords](#).

Valid syntax for custom signature fields

Field	Valid Characters	Usage
HEADER	F-SBID	The header for an attack definition signature. Each custom signature must begin with this header.

Field	Valid Characters	Usage
KEYWORD	<p>Each keyword must start with a pair of dashes (--), and consist of a string of 1 to 19 characters.</p> <p>Normally, keywords are an English word or English words connected by an underscore (_). Keywords are case insensitive.</p>	The keyword is used to identify a parameter.
VALUE	<p>Double quotes (") must be used around the value if it contains a space and/or a semicolon (;).</p> <p>If the value is NULL, the space between the KEYWORD and VALUE can be omitted.</p> <p>Values are case sensitive.</p> <p>Note: If double quotes are used for quoting the value, the double quotes are not considered as part of the value string.</p>	The value is set specifically for a parameter identified by a keyword.

Custom signature keywords

- information
- session
- content
- IP header
- TCP header
- UDP header
- ICMP
- other

Information keywords

attack_id

Syntax: --attack_id <id_int>;

Description:

Use this optional value to identify the signature. It cannot be the same value as any other custom rules. If an attack ID is not specified, the FortiGate automatically assigns an attack ID to the signature. If you are using VDOMs, custom signatures appear only in the VDOM in which you create them. You can use the same attack ID for signatures in different VDOMs.

An attack ID you assign must be between 1000 and 9999.

Example: --attack_id 1234;

name

Syntax: `--name <name_str>;`

Description:

Enter the name of the rule. A rule name must be unique. If you are using VDOMs, custom signatures appear only in the VDOM in which you create them. You can use the same rule name for signatures in different VDOMs. The name you assign must be a string greater than 0 and less than 64 characters in length.

Example: `--name "Buffer_Overflow";`

Session keywords

flow

Syntax: `--flow {from_client[,reversed] | from_server[,reversed] | bi_direction};`

Description:

Specify the traffic direction and state to be inspected. They can be used for all IP traffic.

Example: `--src_port 41523; --flow bi_direction;`

The signature checks traffic to and from port 41523.

If you enable “quarantine attacker”, the optional reversed keyword allows you to change the side of the connection to be quarantined when the signature is detected.

For example, a custom signature written to detect a brute-force log in attack is triggered when “Login Failed” is detected from_server more than 10 times in 5 seconds. If the attacker is quarantined, it is the server that is quarantined in this instance. Adding reversed corrects this problem and quarantines the actual attacker.

Previous FortiOS versions used to_client and to_server values. These are now deprecated, but still function for backwards compatibility.

service

Syntax: `--service {HTTP | TELNET | FTP | DNS | SMTP | POP3 | IMAP | SNMP | RADIUS | LDAP | MSSQL | RPC | SIP | H323 | NBSS | DCERPC | SSH | SSL};`

Description:

Specify the protocol type to be inspected. This keyword allows you to specify the traffic type by protocol rather than by port. If the decoder has the capability to identify the protocol on any port, the signature can be used to detect the attack no matter what port the service is running on. Currently, HTTP, SIP, SSL, and SSH protocols can be identified on any port based on the content.

app_cat

Syntax: `--app_cat <category_int>;`

Description:

Specify the category of the application signature. Signatures with this keyword are considered as application rules. These signatures will appear under Application Control instead of IPS configuration. To display a complete list of application signature categories, enter the following CLI commands:

```
config application list
edit default
config entries
edit 1
set category ?
```

weight

Syntax: `--weight <weight_int>;`

Description:

Specify the weight to be assigned to the signature. This keyword allows a signature with the higher weight to have priority over a signature with a lower weight. This is useful to prioritize between custom and stock signatures and also between different custom signatures.

The weight must be between 0 and 255. Most of the signatures in the Application Control signature database have weights of 10; botnet signatures are set to 250. A range of 20 to 50 is recommended for custom signatures.

Content keywords

byte_extract

Syntax: `byte_extract:<bytes_to_extract>, <offset>, <name> \ [, relative][, multiplier <multiplier value>][, <endian>]\ [, string][, hex][, dec][, oct][, align <align value>][, dce];`

Description:

Use the `byte_extract` option to write rules against length-encoded protocols. This reads some of the bytes from the packet payload and saves it to a variable.

byte_jump

Syntax: `--byte_jump <bytes_to_convert>, <offset>[, multiplier][, relative] [, big] [, little] [, string] [, hex] [, dec] [, oct] [, align];`

Description:

Use the `byte_jump` option to extract a number of bytes from a packet, convert them to their numeric representation, and jump the match reference up that many bytes (for further pattern matching or byte testing). This keyword allows relative pattern matches to take into account numerical values found in network data. The available keyword options include:

- `<bytes_to_convert>`: The number of bytes to examine from the packet.
- `<offset>`: The number of bytes into the payload to start processing.
- `[multiplier]`: multiplier is optional. It must be a numerical value when present. The converted value multiplied by the number is the result to be skipped.
- `relative`: Use an offset relative to last pattern match.
- `big`: Process the data as big endian (default).

- `little`: Process the data as little endian.
- `string`: The data is a string in the packet.
- `hex`: The converted string data is represented in hexadecimal notation.
- `dec`: The converted string data is represented in decimal notation.
- `oct`: The converted string data is represented in octal notation.
- `align`: Round up the number of converted bytes to the next 32-bit boundary.

byte_test

Syntax: `--byte_test <bytes_to_convert>, <operator>, <value>, <offset> [multiplier] [, relative] [, big] [, little] [, string] [, hex] [, dec] [, oct];`

Description:

Use the `byte_test` keyword to compare a byte field against a specific value (with operator). This keyword is capable of testing binary values or converting representative byte strings to their binary equivalent and testing them. The available keyword options include:

- `<bytes_to_convert>`: The number of bytes to compare.
- `<operator>`: The operation to perform when comparing the value (`<`, `>`, `=`, `!`, `&`).
- `<value>`: The value to compare the converted value against.
- `<offset>`: The number of bytes into the payload to start processing.
- `[multiplier]`: multiplier is optional. It must be a numerical value when present. The converted value multiplied by the number is the result to be skipped.
- `relative`: Use an offset relative to last pattern match.
- `big`: Process the data as big endian (default).
- `little`: Process the data as little endian.
- `string`: The data is a string in the packet.
- `hex`: The converted string data is represented in hexadecimal notation.
- `dec`: The converted string data is represented in decimal notation.
- `oct`: The converted string data is represented in octal notation.

depth

Syntax: `--depth <depth_int>;`

Description:

Use the `depth` keyword to search for the contents within the specified number of bytes after the starting point defined by the `offset` keyword. If no offset is specified, the offset is assumed to be equal to 0.

If the value of the `depth` keyword is smaller than the length of the value of the `content` keyword, this signature will never be matched.

The depth must be between 0 and 65535.

distance

Syntax: `--distance <dist_int>;`

Description:

Use the distance keyword to search for the contents within the specified number of bytes relative to the end of the previously matched contents. If the within keyword is not specified, continue looking for a match until the end of the payload.

The distance must be between 0 and 65535.

content

Syntax: `--content [!] "<content_str>";`

Description:

Deprecated, see pattern and context keywords. Use the content keyword to search for the content string in the packet payload. The content string must be enclosed in double quotes.

To have the FortiGate search for a packet that does not contain the specified context string, add an exclamation mark (!) before the content string.

Multiple content items can be specified in one rule. The value can contain mixed text and binary data. The binary data is generally enclosed within the pipe (|) character.

The double quote ("), pipe sign(|) and colon(:) characters must be escaped using a back slash if specified in a content string.

If the value of the content keyword is greater than the length of the value of the depth keyword, this signature will never be matched.

context

Syntax: `--context {uri | header | body | host};`

Description:

Specify the protocol field to look for the pattern. If context is not specified for a pattern, the FortiGate unit searches for the pattern anywhere in the packet buffer. The available context variables are:

- **uri:** Search for the pattern in the HTTP URI line.
- **header:** Search for the pattern in HTTP header lines or SMTP/POP3/SMTP control messages.
- **body:** Search for the pattern in HTTP body or SMTP/POP3/SMTP email body.
- **host:** Search for the pattern in HTTP HOST line.

no_case

Syntax: `--no_case;`

Description:

Use the no-case keyword to force the FortiGate unit to perform a case-insensitive pattern match.

offset

Syntax: `--offset <offset_int>;`

Description:

Use the offset keyword to look for the contents after the specified number of bytes into the payload. The specified number of bytes is an absolute value in the payload. Follow the offset keyword with the depth

keyword to stop looking for a match after a specified number of bytes. If no depth is specified, the FortiGate unit continues looking for a match until the end of the payload.

The offset must be between 0 and 65535.

pattern

Syntax: `--pattern [!]"<pattern_str>"`;

Description:

The FortiGate unit will search for the specified pattern. A pattern keyword normally is followed by a context keyword to define where to look for the pattern in the packet. If a context keyword is not present, the FortiGate unit looks for the pattern anywhere in the packet buffer. To have the FortiGate search for a packet that does not contain the specified URI, add an exclamation mark (!) before the URI.

Example: `--pattern "/level/" --pattern "|E8 D9FF FFFF|/bin/sh" --pattern
!"|20|RTSP/"`

pcre

Syntax: `--pcre [!]"<regex>/[ismxAEGRUB]"`;

Description:

Similarly to the pattern keyword, use the pcre keyword to specify a pattern using Perl-compatible regular expressions (PCRE). A pcre keyword can be followed by a context keyword to define where to look for the pattern in the packet. If no context keyword is present, the FortiGate unit looks for the pattern anywhere in the packet buffer.

For more information about PCRE syntax, go to <http://www.pcre.org>.

The switches include:

- **i:** Case insensitive.
- **s:** Include newlines in the dot metacharacter.
- **m:** By default, the string is treated as one big line of characters. **^** and **\$** match at the beginning and ending of the string. When **m** is set, **^** and **\$** match immediately following or immediately before any newline in the buffer, as well as the very start and very end of the buffer.
- **x:** White space data characters in the pattern are ignored except when escaped or inside a character class.
- **A:** The pattern must match only at the start of the buffer (same as **^**).
- **E:** Set **\$** to match only at the end of the subject string. Without **E**, **\$** also matches immediately before the final character if it is a newline (but not before any other newlines).
- **G:** Invert the "greediness" of the quantifiers so that they are not greedy by default, but become greedy if followed by **?**.
- **R:** Match relative to the end of the last pattern match. (Similar to `distance:0`).
- **U:** Deprecated, see the context keyword. Match the decoded URI buffers.

uri

Syntax: `--uri [!]"<uri_str>"`;

Description:

Deprecated, see pattern and context keywords. Use the uri keyword to search for the URI in the packet payload. The URI must be enclosed in double quotes ("). To have the FortiGate unit search for a packet that does not contain the specified URI, add an exclamation mark (!) before the URI. Multiple content items can be specified in one rule. The value can contain mixed text and binary data. The binary data is generally enclosed within the pipe (|) character. The double quote ("), pipe sign (|) and colon (:) characters must be escaped using a back slash (\) if specified in a URI string.

within

Syntax: `--within <within_int>;`

Description:

Use this together with the distance keyword to search for the contents within the specified number of bytes of the payload.

The within value must be between 0 and 65535.

IP header keywords

dst_addr

Syntax: `--dst_addr [!]<ipv4>;`

Description:

Use the dst_addr keyword to search for the destination IP address. To have the FortiGate search for a packet that does not contain the specified address, add an exclamation mark (!) before the IP address. You can define up to 28 IP addresses or CIDR blocks. Enclose the comma separated list in square brackets.

Example: `dst_addr [172.20.0.0/16, 10.1.0.0/16, 192.168.0.0/16]`

ip_dscp

Syntax: `--ip_dscp`

Description:

Use the ip_dscp keyword to check the IP DSCP field for the specified value.

ip_id

Syntax: `--ip_id <field_int>;`

Description:

Check the IP ID field for the specified value.

ip_option

Syntax: `--ip_option {rr | eol | nop | ts | sec | lsrr | ssrr | satid | any};`

Description:

Use the ip_option keyword to check various IP option settings.

The available options include:

- `rr`: Check if IP RR (record route) option is present.
- `eol`: Check if IP EOL (end of list) option is present.
- `nop`: Check if IP NOP (no op) option is present.
- `ts`: Check if IP TS (time stamp) option is present.
- `sec`: Check if IP SEC (IP security) option is present.
- `lsrr`: Check if IP LSRR (loose source routing) option is present.
- `ssrr`: Check if IP SSRR (strict source routing) option is present.
- `satid`: Check if IP SATID (stream identifier) option is present.
- `any`: Check if IP any option is present.

`ip_tos`

Syntax: `--ip_tos <field_int>;`

Description:

Check the IP TOS field for the specified value.

`ip_ttl`

Syntax: `--ip_ttl [< | >] <ttl_int>;`

Description:

Check the IP time-to-live value against the specified value. Optionally, you can check for an IP time-to-live greater-than (>) or less-than (<) the specified value with the appropriate symbol.

`protocol`

Syntax: `--protocol {<protocol_int> | tcp | udp | icmp};`

Description:

Check the IP protocol header.

Example: `--protocol tcp;`

`src_addr`

Syntax: `--src_addr [!]<ipv4>;`

Description:

Use the `src_addr` keyword to search for the source IP address. To have the FortiGate unit search for a packet that does not contain the specified address, add an exclamation mark (!) before the IP address. You can define up to 28 IP addresses or CIDR blocks. Enclose the comma separated list in square brackets.

Example: `src_addr 192.168.13.0/24`

TCP header keywords

`ack`

Syntax: `--ack <ack_int>;`

Description:

Check for the specified TCP acknowledge number.

dst_port

Syntax: `--dst_port [!]{<port_int> | :<port_int> | <port_int>: | <port_int>:<port_int>;}`

Description:

Use the `dst_port` keyword to specify the destination port number.

You can specify a single port or port range:

- `<port_int>` is a single port.
- `:<port_int>` includes the specified port and all lower numbered ports.
- `<port_int>:` includes the specified port and all higher numbered ports.
- `<port_int>:<port_int>` includes the two specified ports and all ports in between.

seq

Syntax: `--seq [operator,]<number>[,relative];`

Description:

Check for the specified TCP sequence number.

- `operator` includes `=,<,>,!`.
- `relative` indicates it's relative to the initial sequence number of the TCP session.

src_port

Syntax: `--src_port [!]{<port_int> | :<port_int> | <port_int>: | <port_int>:<port_int>;}`

Description:

Use the `src_port` keyword to specify the source port number. You can specify a single port or port range:

- `<port_int>` is a single port.
- `:<port_int>` includes the specified port and all lower numbered ports.
- `<port_int>:` includes the specified port and all higher numbered ports.
- `<port_int>:<port_int>` includes the two specified ports and all ports in between.

tcp_flags

Syntax: `--tcp_flags <SAFRUP120>[!|*|+] [,<SAFRUP120>];`

Description:

Specify the TCP flags to match in a packet.

- S: Match the SYN flag.
- A: Match the ACK flag.
- F: Match the FIN flag.
- R: Match the RST flag.

- U: Match the URG flag.
- P: Match the PSH flag.
- 1: Match Reserved bit 1.
- 2: Match Reserved bit 2.
- 0: Match No TCP flags set.
- !: Match if the specified bits are not set.
- *: Match if any of the specified bits are set.
- +: Match on the specified bits, plus any others.

The first part of the value (<SAFRUP120>) defines the bits that must be present for a successful match.

Example:

`--tcp_flags AP` only matches the case where both A and P bits are set.

The second part ([, <SAFRUP120>]) is optional, and defines the additional bits that can be present for a match.

For example `tcp_flags S,12` matches the following combinations of flags: S, S and 1, S and 2, S and 1 and 2. The modifiers !, * and + cannot be used in the second part.

window_size

Syntax: `--window_size [!]<window_int>;`

Description:

Check for the specified TCP window size. You can specify the window size as a hexadecimal or decimal integer. A hexadecimal value must be preceded by 0x. To have the FortiGate search for the absence of the specified window size, add an exclamation mark (!) before the window size.

UDP header keywords

dst_port

Syntax: `--dst_port [!]{<port_int> | :<port_int> | <port_int>: | <port_int>:<port_int>;}`

Description:

Specify the destination port number. You can specify a single port or port range:

- <port_int> is a single port.
- :<port_int> includes the specified port and all lower numbered ports.
- <port_int>: includes the specified port and all higher numbered ports.
- <port_int>:<port_int> includes the two specified ports and all ports in between.

src_port

Syntax: `--src_port [!]{<port_int> | :<port_int> | <port_int>: | <port_int>:<port_int>;}`

Description:

Specify the destination port number. You can specify a single port or port range:

- <port_int> is a single port.
- :<port_int> includes the specified port and all lower numbered ports.
- <port_int>: includes the specified port and all higher numbered ports.
- <port_int>:<port_int> includes the two specified ports and all ports in between.

ICMP keywords

icmp_code

Syntax: --icmp_code <code_int>;

Description:

Specify the ICMP code to match.

icmp_id

Syntax: --icmp_id <id_int>;

Description:

Check for the specified ICMP ID value.

icmp_seq

Syntax: --icmp_seq <seq_int>;

Description:

Check for the specified ICMP sequence value.

icmp_type

Syntax: --icmp_type <type_int>;

Description:

Specify the ICMP type to match.

Other keywords

data_size

Syntax: --data_size {<size_int> | <<size_int> | >>size_int};

Description:

Test the packet payload size. With data_size specified, packet reassembly is turned off automatically. So a signature with data_size and only_stream values set is wrong.

- <size_int> is a particular packet size.
- <<size_int> is a packet smaller than the specified size.
- >>size_int> is a packet larger than the specified size.

Examples:

- `--data_size 300;`
- `--data_size <300;`
- `--data_size >300;`

data_at

Syntax: `--data_at <offset_int>[, relative];`

Description:

Verify that the payload has data at a specified offset, optionally looking for data relative to the end of the previous content match.

dump-all-html

Syntax: `--dump-all-html`

Description:

Dump all HTML files for benchmarking via iSniff. When there is no file type specified, all HTML files are dumped.

rate

Syntax: `--rate <matches_int>,<time_int>;`

Description:

Instead of generating log entries every time the signature is detected, use this keyword to generate a log entry only if the signature is detected a specified number of times within a specified time period.

- `<matches_int>` is the number of times a signature must be detected.
- `<time_int>` is the length of time in which the signature must be detected, in seconds.

For example, if a custom signature detects a pattern, a log entry will be created every time the signature is detected. If `--rate 100,10;` is added to the signature, a log entry will be created if the signature is detected 100 times in the previous 10 seconds. Use this command with `--track` to further limit log entries to when the specified number of detections occur within a certain time period involving the same source or destination address rather than all addresses.

rpc_num

Syntax: `--rpc_num <app_int>[, <ver_int> | *][, <proc_int> | *];`

Description:

Check for RPC application, version, and procedure numbers in SUNRPC CALL requests. The * wild card can be used for version and procedure numbers.

same_ip

Syntax: `--same_ip;`

Description:

Check that the source and the destination have the same IP addresses.

track

Syntax: `--track {SRC_IP | DST_IP | DHCP_CLIENT | DNS_DOMAIN}[,block_int];`

Description:

When used with `--rate`, this keyword narrows the custom signature rate totals to individual addresses.

- `SRC_IP`: tracks the packet's source IP.
- `DST_IP`: tracks the packet's destination IP.
- `DHCP_CLIENT`: tracks the DHCP client's MAC address.
- `DNS_DOMAIN`: counts the number of any specific domain name.
- `block_int` has the FortiGate unit block connections for the specified number of seconds, from the client or to the server, depending on which is specified.

For example, if `--rate 100,10` is added to the signature, a log entry will be created if the signature is detected 100 times in the previous 10 seconds. The FortiGate unit maintains a single total, regardless of source and destination address.

If the same custom signature also includes `--track client`; matches are totaled separately for each source address. A log entry is added when the signature is detected 100 times in 10 seconds within traffic from the same source address.

The `--track` keyword can also be used without `--rate`. If an integer is specified, the client or server will be blocked for the specified number of seconds every time the signature is detected.

Creating a custom signature to block access to example.com

In this first example, you will create a custom signature to block access to the example.com URL.

This example describes the use of the custom signature syntax to block access to a URL. To create the custom signature entry in the FortiGate's GUI, see [Custom Application & IPS Signatures](#).

1. Enter the custom signature basic format.

All custom signatures have a header and at least one keyword/value pair. The header is always the same:

```
F-SBID( )
```

The keyword/value pairs appear within the parentheses and each pair is followed by a semicolon.

2. Choose a name for the custom signature

Every custom signature requires a name, so it is a good practice to assign a name before adding any other keywords. Use the `--name` keyword to assign the custom signature a name. The name value follows the keyword after a space. Enclose the name value in double-quotes:

```
F-SBID( --name "Block.example.com"; )
```

The signature, as it appears here, will not do anything if you try to use it. It has a name, but does not look for any patterns in network traffic. You must specify a pattern that the FortiGate unit will search for.

3. Add a signature pattern

Use the `--pattern` keyword to specify what the FortiGate unit will search for:

```
F-SBID( --name "Block.example.com"; --pattern "example.com"; )
```

The signature will now detect the example.com URL appearing in network traffic. The custom

signature should only detect the URL in HTTP traffic, however. Any other traffic with the URL should be allowed to pass. For example, an email message to or from example.com should not be stopped.

4. Specify the service

Use the `--service` keyword to limit the effect of the custom signature to only the HTTP protocol.

```
F-SBID( --name "Block.example.com"; --pattern "example.com"; --service HTTP; )
```

The FortiGate unit will limit its search for the pattern to the HTTP protocol. Even though the HTTP protocol uses only TCP traffic, the FortiGate will search for HTTP protocol communication in TCP, UDP, and ICMP traffic. This is a waste of system resources that you can avoid by limiting the search further, as shown below.

5. Specify the traffic type.

Use the `--protocol tcp` keyword to limit the effect of the custom signature to only TCP traffic. This will save system resources by not unnecessarily scanning UDP and ICMP traffic.

```
F-SBID( --name "Block.example.com"; --pattern "example.com"; --service HTTP; --
protocol tcp; )
```

The FortiGate unit will limit its search for the pattern to TCP traffic and ignore UDP and ICMP network traffic.

6. Ignore case sensitivity

By default, patterns are case sensitive. If a user directed his or her browser to Example.com, the custom signature would not recognize the URL as a match.

Use the `--no_case` keyword to make the pattern matching case insensitive.

```
F-SBID( --name "Block.example.com"; --pattern "example.com"; --service HTTP; --
protocol tcp; --no_case; )
```

Unlike all of the other keywords in this example, the `--no_case` keyword has no value. Only the keyword is required.

7. Limit pattern scans to only traffic sent from the client

The `--flow` command can be used to further limit the network traffic being scanned to only that sent by the client or by the server.

```
F-SBID( --name "Block.example.com"; --pattern "example.com"; --service HTTP; --
protocol tcp; --no_case; --flow from_client; )
```

Web servers do not contact clients until clients first open a communication session. Therefore, using the `--flow from_client` command will force the FortiGate to ignore all traffic from the server. Since the majority of HTTP traffic flows from the server to the client, this will save considerable system resources and still maintain protection.

8. Specify the context

When the client browser tries to contact example.com, a DNS is first consulted to get the example.com server IP address. The IP address is then specified in the URL field of the HTTP communication. The domain name will still appear in the host field, so this custom signature will not function without the `--context host` keyword/value pair.

```
F-SBID( --name "Block.example.com"; --pattern "example.com"; --service HTTP; --no_
case; --flow from_client; --context host; )
```

Creating a custom signature to block the SMTP “vrfy” command

The SMTP “vrfy” command can be used to verify the existence of a single email address or to list all of the valid email accounts on an email server. A spammer could potentially use this command to obtain a list of all valid email users and direct spam to their inboxes.

In this example, you will create a custom signature to block the use of the vrfy command. Since the custom signature blocks the vrfy command from coming through the FortiGate unit, the administrator can still use the command on the internal network.

This example describes the use of the custom signature syntax to block the vrfy command. To create the custom signature entry in the FortiGate's GUI, see [Custom Application & IPS Signatures](#).

1. Enter the custom signature basic format

All custom signatures have a header and at least one keyword/value pair. The header is always the same:

```
F-SBID( )
```

The keyword/value pairs appear within the parentheses and each pair is followed by a semicolon.

2. Choose a name for the custom signature

Every custom signature requires a name, so it is a good practice to assign a name before you add any other keywords.

Use the `--name` keyword to assign the custom signature a name. The name value follows the keyword after a space. Enclose the name value in double-quotes:

```
F-SBID( --name "Block.SMTP.VRFY.CMD"; )
```

The signature, as it appears here, will not do anything if you try to use it. It has a name, but does not look for any patterns in network traffic. You must specify a pattern that the FortiGate unit will search for.

3. Add a signature pattern

Use the `--pattern` keyword to specify what the FortiGate unit will search for:

```
F-SBID( --name "Block.SMTP.VRFY.CMD"; --pattern "vrfy"; )
```

The signature will now detect the vrfy command appearing in network traffic. The custom signature should only detect the command in SMTP traffic, however. Any other traffic with the pattern should be allowed to pass. For example, an email message discussing the vrfy command should not be stopped.

4. Specify the service.

Use the `--service` keyword to limit the effect of the custom signature to only the HTTP protocol.

```
F-SBID( --name "Block.SMTP.VRFY.CMD"; --pattern "vrfy"; --service SMTP; )
```

The FortiGate unit will limit its search for the pattern to the SMTP protocol.

Even though the SMTP protocol uses only TCP traffic, the FortiGate will search for SMTP protocol communication in TCP, UDP, and ICMP traffic. This is a waste of system resources that you can avoid by limiting the search further, as shown below.

5. Specify the traffic type.

Use the `--protocol tcp` keyword to limit the effect of the custom signature to only TCP traffic. This will save system resources by not unnecessarily scanning UDP and ICMP traffic.

```
F-SBID( --name "Block.SMTP.VRFY.CMD"; --pattern "vrfy"; --service SMTP; --
        protocol tcp; )
```

The FortiGate unit will limit its search for the pattern to TCP traffic and ignore the pattern in UDP and ICMP network traffic.

6. Ignore case sensitivity.

By default, patterns are case sensitive. If a user directed his or her browser to Example.com, the custom signature would not recognize the URL as a match.

Use the `--no_case` keyword to make the pattern matching case insensitive.

```
F-SBID( --name "Block.SMTP.VRFY.CMD"; --pattern "vrfy"; --service SMTP; --no_case; )
```

Unlike all of the other keywords in this example, the `--no_case` keyword has no value. Only the keyword is required.

7. Specify the context.

The `SMTP vrfy` command will appear in the SMTP header. The `--context host keyword/value` pair allows you to limit the pattern search to only the header.

```
F-SBID( --name "Block.SMTP.VRFY.CMD"; --pattern "vrfy"; --service SMTP; --no_case; -
        -context header; )
```

Creating a custom signature to block files according to the file's hash value

In this example, you will create a custom signature that allows you to specify a hash value (or checksum) of a file that you want to block. To block multiple files you can create a custom signature for each file with that file's hash value in it and then add all of the custom signatures to an IPS sensor and set the action to block for each one. When IPS encounters a file with a matching hash value the file is blocked.

This example uses a CRC32 checksum of the file as the hash value of the file to be blocked. You can use any utility that supports CRC32 checksums to generate the hash value.

1. Enter the custom signature basic format.

All custom signatures have a header and at least one keyword/value pair. The header is always the same:

```
F-SBID( )
```

The keyword/value pairs appear within the parentheses and each pair is followed by a semicolon.

2. Choose a name for the custom signature

Every custom signature requires a name, so it is a good practice to assign a name before adding any other keywords. Use the `--name` keyword to assign the custom signature a name. The name value follows the keyword after a space. Enclose the name value in double-quotes:

```
F-SBID( --name "File.Hash.Example"; )
```

The signature, as it appears here, will not do anything if you try to use it. It has a name, but does not look for any patterns in network traffic.

3. Specify the traffic type.

Use the `--protocol tcp` keyword to limit the effect of the custom signature to only TCP traffic. This will save system resources by not unnecessarily scanning UDP and ICMP traffic.

```
F-SBID( --name "File.Hash.Example"; --protocol tcp; )
```

The FortiGate unit will limit its search for the pattern to TCP traffic and ignore UDP and ICMP network traffic.

4. Add the CRC32 hash value.

Use the `--crc32` keyword. This indicates that the value that follows is a hexadecimal number that represents the CRC32 checksum of the file. The `--crc32` keyword also requires that you include the file length. The syntax is `--crc32 <checksum>,<file-length>;`. The following example shows the syntax for a file with checksum 51480492 and file length 822.

```
F-SBID( --name "File.Hash.Example"; --protocol tcp; --crc32 51480492,822; )
```

Other security profiles considerations

The following topics are included in this section:

- [Security profiles and Virtual Domains \(VDOMs\)](#)
- [Conserve mode](#)
- [Using wildcards and Perl regular expressions](#)
- [CPU allocation and tuning commands to survive reboot](#)

Global security profiles across Virtual domains (VDOMs)

Previously, if you enabled virtual domains (VDOMs) on your FortiGate unit, any Security Profiles configuration was limited to the VDOM in which you configured it.

Now Security Profiles can be configured globally across multiple VDOMs. In many VDOM environments, some or all profiles may be commonly-shared, for example an MSSP with "parental controls" configured will most likely have the same Web Filtering and Application Control profiles per VDOM.

Global profiles are configured under **Global > Security Profiles** in the GUI or under the following `config global` commands in the CLI:

- `antivirus profile`
- `application list`
- `dlp sensor`
- `ips sensor`
- `webfilter profile`

The name for any global profile must start with "g-" for identification. Global profiles are available as read-only for VDOM-level administrators and can only be edited or deleted from within the global settings.

Each security feature has at least one default global profile, available for all VDOMs.

Both Global security profile configuration and the various databases used by Security Profiles features are shared. The FortiGuard antivirus and IPS databases and updates to the databases are shared. The FortiGuard web filter and spam filter features access the FortiGuard distribution network and read the same information when checking email for spam and web site categories and classification.

Conserve mode

FortiGate units perform all Security Profiles processing in physical RAM. Since each model has a limited amount of memory, conserve mode is activated when the remaining free memory is nearly exhausted or the AV proxy has reached the maximum number of sessions it can service. While conserve mode is active, the AV proxy does not accept new sessions.

A warning will appear in the top bar of the FortiGate, regardless of which page in the FortiGate GUI you are on.

The AV proxy

Most content inspection the FortiGate unit performs requires that the files, email messages, URLs, and web pages be buffered and examined as a whole. The AV proxy performs this function, and because it may be buffering many files at the same time, it uses a significant amount of memory. Conserve mode is designed to

prevent all the component features of the FortiGate unit from trying to use more memory than it has. Because the AV proxy uses so much memory, conserve mode effectively disables it in most circumstances. As a result, the content inspection features that use the AV proxy are also disabled in conserve mode.

All of the Security Profiles features use the AV proxy with the exception of IPS, application control, DoS as well as flow-based antivirus, DLP, and web filter scanning. These features continue to operate normally when the FortiGate unit enters conserve mode.

Entering and exiting conserve mode

A FortiGate unit will enter conserve mode because it is nearly out of physical memory, or because the AV proxy has reached the maximum number of sessions it can service. The memory threshold that triggers conserve mode varies by model, but it is about 20% free memory. When memory use rises to the point where less than 20% of the physical memory is free, the FortiGate unit enters conserve mode.

The FortiGate unit will leave conserve mode only when the available physical memory exceeds about 30%. When exiting conserve mode, all new sessions configured to be scanned with features requiring the AV proxy will be scanned as normal, with the exception of a unit configured with the one-shot option.

Conserve mode effects

What happens when the FortiGate unit enters conserve mode depends on how you have `av-failopen` configured. There are four options:

off

The off setting forces the FortiGate unit to stop all traffic that is configured for content inspection by Security Profiles features that use the AV proxy. New sessions are not allowed but current sessions continue to be processed normally unless they request more memory. Sessions requesting more memory are terminated.

For example, if a security policy is configured to use antivirus scanning, the traffic it permits is blocked while in conserve mode. A policy with IPS scanning enabled continues as normal. A policy with both IPS and antivirus scanning is blocked because antivirus scanning requires the AV proxy.

Use the off setting when security is more important than a loss of access while the problem is rectified.

pass

The pass setting allows traffic to bypass the AV proxy and continue to its destination. Since the traffic is bypassing the proxy, no Security Profiles scanning that requires the AV proxy is performed. Security Profiles scanning that does not require the AV proxy continues normally.

Use the pass setting when access is more important than security while the problem is rectified.

Pass is the default setting.

one-shot

The one-shot setting is similar to pass in that traffic is allowed when conserve mode is active. The difference is that a system configured for one-shot will force new sessions to bypass the AV proxy even after it leaves conserve mode. The FortiGate unit resumes use of the AV proxy only when the `av-failopen` setting is changed or the unit is restarted.

idledrop

The idledrop setting will recover memory and session space by terminating all the sessions associated with the host that has the most sessions open. The FortiGate may force this session termination a number of times, until enough memory is available to allow it to leave conserve mode.

The idledrop setting is primarily designed for situations in which malware may continue to open sessions until the AV proxy cannot accept more new sessions, triggering conserve mode. If your FortiGate unit is operating near capacity, this setting could cause the termination of valid sessions. Use this option with caution.

Configuring the av-failopen command

You can configure the av-failopen command using the CLI.

```
config system global
    set av-failopen {off | pass | one-shot | idledrop}
end
```

The default setting is pass.

Using wildcards and Perl regular expressions

Many Security Profiles feature list entries can include wildcards or Perl regular expressions.

For more information about using Perl regular expressions, see <http://perldoc.perl.org/perlretut.html>.

Regular expression vs. wildcard match pattern

A wildcard character is a special character that represents one or more other characters. The most commonly used wildcard characters are the asterisk (*), which typically represents zero or more characters in a string of characters, and the question mark (?), which typically represents any one character.

In Perl regular expressions, the '.' character refers to any single character. It is similar to the '?' character in wildcard match pattern. As a result:

- example.com not only matches example.com but also examplea.com, exampleb.com, examplec.com, and so on.



To add a question mark (?) character to a regular expression from the FortiGate CLI, enter Ctrl+V followed by ?. To add a single backslash character (\) to a regular expression from the CLI you must add precede it with another backslash character. For example, `example\\.com`.

To match a special character such as '.' and '*' use the escape character '\\'. For example:

- To match example.com, the regular expression should be: `example\\.com`

In Perl regular expressions, '*' means match 0 or more times of the character before it, not 0 or more times of any character. For example:

- `exam*.com` matches exammmm.com but does not match example.com

To match any character 0 or more times, use '.*' where '.' means any character and the '*' means 0 or more times. For example, the wildcard match pattern `exam*.com` should be `exam.*\\.com`.

Word boundary

In Perl regular expressions, the pattern does not have an implicit word boundary. For example, the regular expression “test” not only matches the word “test” but also any word that contains “test” such as “atest”, “mytest”, “testimony”, “atestb”. The notation “\b” specifies the word boundary. To match exactly the word “test”, the expression should be \btest\b.

Case sensitivity

Regular expression pattern matching is case sensitive in the web and Email Filter filters. To make a word or phrase case insensitive, use the regular expression /i. For example, /bad language/i will block all instances of “bad language”, regardless of case.

Perl regular expression formats

The following table lists and describes some example Perl regular expressions.

Perl regular expression formats

Expression	Matches
abc	“abc” (the exact character sequence, but anywhere in the string)
^abc	“abc” at the beginning of the string
abc\$	“abc” at the end of the string
a b	Either “a” or “b”
^abc abc\$	The string “abc” at the beginning or at the end of the string
ab{2,4}c	“a” followed by two, three or four “b”s followed by a “c”
ab{2,}c	“a” followed by at least two “b”s followed by a “c”
ab*c	“a” followed by any number (zero or more) of “b”s followed by a “c”
ab+c	“a” followed by one or more b's followed by a c
ab?c	“a” followed by an optional “b” followed by a “c”; that is, either “abc” or “ac”
a.c	“a” followed by any single character (not newline) followed by a “c”
a\.c	“a.c” exactly
[abc]	Any one of “a”, “b” and “c”
[Aa]bc	Either of “Abc” and “abc”
[abc]+	Any (nonempty) string of “a”s, “b”s and “c”s (such as “a”, “abba”, “acbabcacaa”)

Expression	Matches
[^abc]+	Any (nonempty) string which does not contain any of “a”, “b”, and “c” (such as “defg”)
\d\d	Any two decimal digits, such as 42; same as \d{2}
/i	Makes the pattern case insensitive. For example, <code>/bad language/i</code> blocks any instance of <code>bad language</code> regardless of case.
\w+	A “word”: A nonempty sequence of alphanumeric characters and low lines (underscores), such as <code>foo</code> and <code>12bar8</code> and <code>foo_1</code>
100\s*mk	The strings “100” and “mk” optionally separated by any amount of white space (spaces, tabs, newlines)
abc\b	“abc” when followed by a word boundary (for example, in “abc!” but not in “abcd”)
perl\b	“perl” when not followed by a word boundary (for example, in “perlert” but not in “perl stuff”)
\x	Tells the regular expression parser to ignore white space that is neither preceded by a backslash character nor within a character class. Use this to break up a regular expression into (slightly) more readable parts.
/x	Used to add regular expressions within other text. If the first character in a pattern is forward slash '/', the '/' is treated as the delimiter. The pattern must contain a second '/'. The pattern between '/' will be taken as a regular expressions, and anything after the second '/' will be parsed as a list of regular expression options ('i', 'x', etc). An error occurs if the second '/' is missing. In regular expressions, the leading and trailing space is treated as part of the regular expression.

Examples of regular expressions

Block any word in a phrase

```
/block|any|word/
```

Block purposely misspelled words

Spammers often insert other characters between the letters of a word to fool spam blocking software.

```
/^.*v.*i.*a.*g.*r.*o.*$/i
/cr[eéèêë] [\+ \- \* = < > \. \, ; ! \? % & $ @ \^ \° \ $ £ € \{ \} ( ) \[ \] \\\\_01]dit/i
```

Block common spam phrases

The following phrases are some examples of common phrases found in spam messages.

```
/try it for free/i
/student loans/i
/you're already approved/i
/special [\+ \- \* = < > \. \, ; ! \? % & ~ # $ @ \^ \° \ $ £ € \{ \} ( ) \[ \] \\\\_1]offer/i
```

Control how sessions are distributed to Fortinet processes

Previously, the explicit web proxy balanced the client to a specific WAD daemon based only on the source IP. There are cases where customers use another explicit proxy in front of the FortiGate. With such a design, the FortiGate can see the traffic originating from only one IP address (or a small set of IP addresses) and utilize only one (or a small number) of WAD processes.

This new feature modifies the wad-worker balancing algorithm to also use the source port in addition to source IP when distributing the client to a specific WAD daemon. With this in place, even the connections from one IP address will be balanced over all the WAD processes. This also avoids the degraded performance results for the cases where customers are testing the FortiGate as the explicit webproxy to replace Bluecoats, but don't want to remove Bluecoats from the network for the PoC.

Syntax

```
config system global
    set wad-source-affinity {enable | disable}
end
```

This feature is enabled by default. Disabling this option results in some features to be unsupported. IP-based user authentication, disclaimer messages, security profile override, authentication cookies, MAPI scanning, and some video caches such as Youtube are not supported.

CPU allocation and tuning commands to survive reboot

CPU affinity, whereby a process will execute on a specific CPU, can be changed so it survives a reboot.

CLI syntax:

```
config system global
    set av-affinity
    set ips-affinity
    set miglog-affinity
end
```

av-affinity: Affinity setting for AV scanning (64-bit hexadecimal value in the format of xxxxxxxx_xxxxxxxx).

ips-affinity: Affinity setting for IPS (64-bit hexadecimal value in the format of xxxxxxxx_xxxxxxxx; allowed CPUs must be less than total number of IPS engine daemons). This option is only available if the FortiGate includes NP6 processors and support NTurbo.

miglog-affinity: Affinity setting for logging (64-bit hexadecimal value in the format of xxxxxxxx_xxxxxxxx).

Excluding industrial IP signatures

To reduce performance impacts caused by industrial IP signatures, the admin can choose to exclude the industrial signatures when they are loaded by IPS; the industrial signatures then become inactive as a result. The following CLI command has been restored for this purpose.

Syntax

```
config ips global
    set exclude-signatures {none | industrial}
```

end

Chapter 22 - Server Load Balancing

This FortiOS Handbook chapter contains the following sections:

[Inside FortiOS: Server Load Balancing](#) highlights the features and benefits of FortiOS server load balancing.

[Basic load balancing configuration example](#) introduces FortiOS server load balancing by providing a basic configuration example.

[Configuring load balancing](#) describes how to configure FortiOS server load balancing.

[HTTP and HTTPS load balancing, multiplexing, and persistence](#) describes FortiOS server load balancing features that support load balancing HTTP and HTTPS sessions.

[SSL/TLS load balancing](#) describes FortiOS server load balancing features that support load balancing SSL and TLS sessions.

[IP, TCP, and UDP load balancing](#) describes FortiOS server load balancing features that support load balancing IP, TCP, and UDP sessions.

The following configuration examples are also included:

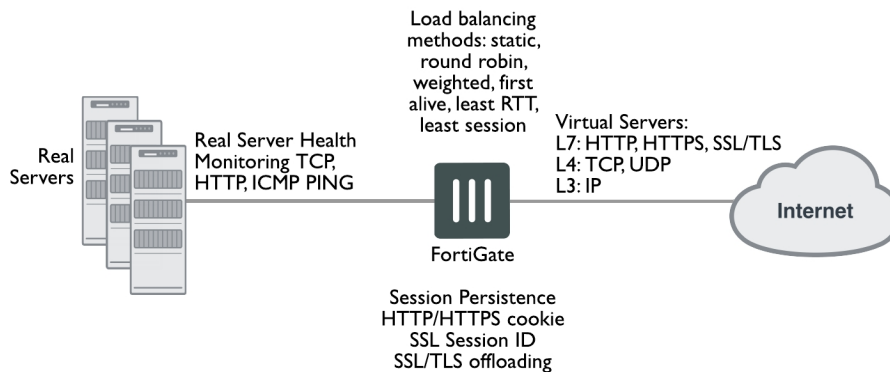
- [Example HTTP load balancing to three real web servers](#)
- [Example Basic IP load balancing configuration](#)
- [Example Adding a server load balance port forwarding virtual IP](#)
- [Example Weighted load balancing configuration](#)
- [Example HTTP and HTTPS persistence configuration](#)

Inside FortiOS: Server Load Balancing

Server load balancing distributes workloads across multiple network servers, allowing simultaneous IPv4, IPv6, IPv4 to IPv6 and IPv6 to IPv4 requests to be handled quickly and reliably.

Server Load Balancing combined with NGFW and UTM protection

By introducing comprehensive server load balancing functionality to Next Generation Firewall (NGFW) and Unified Threat Management (UTM) solutions FortiOS takes threat protection to a whole new level. Rather than going to the expense of deploying multiple solutions to protect your server farm, you can combine firewalling, NGFW, UTM and load balancing into a single FortiGate unit or cluster. The benefit of consolidation is not only limited to cost.



Key Features & Benefits

Increased resilience	A consolidated solution results in significantly simplified network architecture. High availability can be provided for all technologies with just a pair of devices rather than several.
Reduced operational overheads	A unified management solution consisting of a single GUI, logging and reporting, SNMP monitoring and other management functions will significantly reduce the resources required to manage the multiple technology areas. A consolidated solution provides a single point of contact for support and renewals rather than having to deal with multiple vendors.

The FortiOS server load balancing feature set contains all of the features you would expect of a server load balancing solution. Traffic can be balanced across multiple backend servers based on multiple load balancing schedules including static (failover), round robin, weighted to account for different sized servers, or based on the health and performance of the server including round trip time and number of connections.

The load balancer supports HTTP, HTTPS, IMAPS, POP3S, SMTPS, SSL/TLS, and generic TCP/UDP and IP protocols. Session persistence is supported based on the SSL session ID, based on an injected HTTP cookie, or based on the HTTP or HTTPS host. SSL/TLS load balancing includes protection from protocol downgrade attacks. Server load balancing is supported on most FortiGate devices and includes up to 10,000 virtual servers on our high end systems.

SSL/TLS offloading

With more and more critical business applications being made available online and in the cloud, the demand for secure remote continues to increase. While securing web and email applications with SSL/TLS is essential, this protection adds significant performance overheads. An SSL/TLS protected application running on a standard server will perform all the costly encryption/decryption and key exchange routines in software which uses vital CPU resources that should be available for running the application. The consequence of this is that many more or more powerful servers are required to deliver the application.

FortiGate SSL/TLS offloading is designed with the explosion of SSL/TLS applications in mind. The key exchange and encryption/decryption tasks are offloaded to the FortiGate unit where they are accelerated using FortiASIC technology providing significantly more performance than a standard server or load balancer could handle. This frees up valuable resources on the server farm which can be used to run a more responsive business. Server load

balancing offloads most SSL/TLS versions including SSL 3.0, TLS 1.0 and TLS 1.2 and supports full mode or half mode SSL offloading with DH key sizes up to 4096 bits.

SSL/TLS content inspection

Traditionally, SSL encrypted application data would be invisible to any border gateway filtering solution. This is because the encryption process prevents the payload of any connection from being seen other than by the communicating systems. FortiGate SSL Offloading allows the application payload to be inspected before it reaches your servers; preventing intrusion attempts, blocking viruses, stopping unwanted applications, and preventing data leakage. SSL/TLS content inspection supports TLS versions 1.0, 1.1, and 1.2 and SSL versions 1.0, 1.1, 1.2, and 3.0.

Health Check

Health checking can be enabled to prevent load balancing traffic from being sent to a non-functioning real server. Real server health can be monitored using ICMP ping or more sophisticated TCP testing. The most comprehensive test is HTTP which verifies that the HTTP application is responding and that it is returning the correct content.

Health checking removes real servers from the load balancing cluster which are returning invalid content. The removal of real servers from the clusters is based on the Interval, Timeout and Retry Settings:

Interval	How often to test the server.
Timeout	What maximum response time is permissible before a server is treated as non-functional.
Retry	How many failures before the server is considered “dead” and removed from the cluster.

Server Monitoring and Management

The health and performance of real servers can be monitored from the FortiGate GUI. Virtual servers and their assigned real servers can be monitored for health status, if there have been any monitor events, number of active sessions, round trip time and number of bytes processed. Should a server become problematic and require administration, it can be gracefully removed from the Real Server pool to enable disruption free maintenance. When a removed real server is able to operate it can gracefully be added back to the virtual server.

HTTP Multiplexing

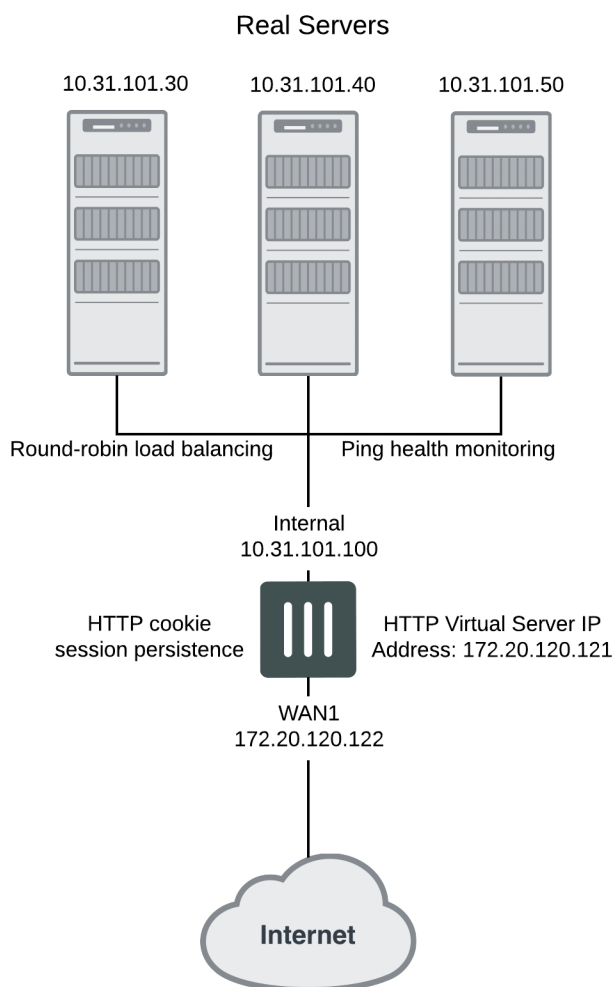
A performance saving feature of HTTP/1.1 compliant web servers is the ability to pipeline requests on the same connection. This allows a single HTTPD process on the server to interleave and server multiple requests. HTTP multiplexing reduces the number idle sessions, too many of which can exhaust the resources on a server. The Fortinet solution has the ability to take multiple separate inbound sessions and multiplex them over the same internal session. This reduces the load on the backend server and increases the overall performance.

Basic load balancing configuration example

This section describes the steps required to configure the load balancing configuration shown below. In this configuration a FortiGate-51B unit is load balancing HTTP traffic from the Internet to three HTTP servers on the Internal network. HTTP sessions are accepted at the wan1 interface with destination IP address 172.20.120.121 on TCP port 8080 and forwarded from the internal interface to the web servers. When forwarded the destination address of the sessions is translated to the IP address of one of the web servers.

The load balancing configuration also includes session persistence using HTTP cookies, round-robin load balancing, and TCP health monitoring for the real servers. Ping health monitoring consists of the FortiGate unit using ICMP ping to make sure the web servers can respond to network traffic.

Virtual server and real servers setup



To configure the example load balancing configuration - general configuration steps

1. Add a load balance ping health check monitor.
A ping health check monitor causes the FortiGate unit to ping the real servers every 10 seconds. If one of the servers does not respond within 2 seconds, the FortiGate unit will retry the ping 3 times before assuming that the HTTP server is not responding.
2. Add a load balance virtual server.
3. Add the three load balance real servers to the virtual server.
4. Add a security policy that includes the load balance virtual server as the destination address.

To configure the example load balancing configuration

1. Go to **Policy & Objects > Health Check** and add the following health check monitor.

Name	Ping-mon-1
Type	Ping
Interval	10 seconds
Timeout	2 seconds
Retry	3

2. Go to **Policy & Objects > Virtual Servers** and add a virtual server that accepts the traffic to be load balanced.

Name	Vserver-HTTP-1
Type	HTTP
Interface	wan1
Virtual Server IP	172.20.120.121
Virtual Server Port	8080
Load Balance Method	Round Robin
Persistence	HTTP Cookie
Health Check	Ping-mon-1
HTTP Multiplexing	Do not select
Preserve Client IP	Do not select

3. On the same GUI page add the real servers to the virtual server.

IP Address	10.31.101.30
Port	80
Max Connections	0
Mode	Active

IP Address	10.31.101.40
Port	80
Max Connections	0
Mode	Active

IP Address	10.31.101.50
Port	80
Max Connections	0
Mode	Active

- Go to **Policy & Objects > IPv4 Policy** and add a wan1 to internal security policy that includes the virtual server. This policy also applies an Antivirus profile to the load balanced sessions.

Name	Example-policy
Incoming Interface	wan1
Outgoing Interface	internal
Source	all
Destination	Vserver-HTTP-1
Schedule	always
Service	ALL
Action	ACCEPT
NAT	Turn on NAT and select Use Outgoing Interface Address .
Antivirus	Turn on and select an Antivirus profile.

- Select **OK**.

To configure the example load balancing configuration from the CLI

- Use the following command to add a Ping health check monitor.

```
config firewall ldb-monitor
edit ping-mon-1
set type ping
set interval 10
set timeout 2
set retry 3
end
```

- Use the following command to add the virtual server that accepts HTTP sessions on port 8080 at the wan1 interface and load balances the traffic to three real servers.

```
config firewall vip
```

```
edit Vserver-HTTP-1
    set type server-load-balance
    set server-type http
    set ldb-method round-robin
    set extip 172.20.120.30
    set extintf wan1
    set extport 8080
    set persistence http-cookie
    set monitor tcp-mon-1
    config realservers
        edit 1
            set ip 10.31.101.30
            set port 80
        next
        edit 2
            set ip 10.31.101.40
            set port 80
        end
        edit 3
            set ip 10.31.101.50
            set port 80
        end
    end
end
```

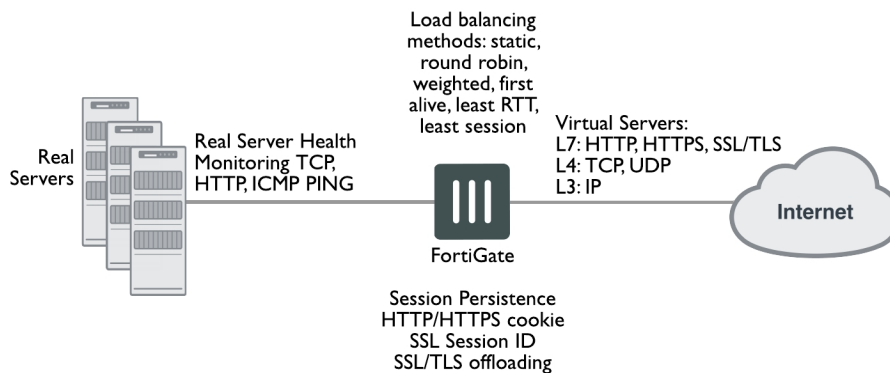
3. Use the following command to add a security policy that includes the load balance virtual server as the destination address.

```
config firewall policy
    edit 0
        set srcintf wan1
        set srcaddr all
        set dstintf internal
        set dstaddr Vserver-HTTP-1
        set action accept
        set schedule always
        set service ALL
        set nat enable
        set utm-status enable
        set profile-protocol-options default
        set av-profile scan
    end
```

Configuring load balancing

This section describes how to use the FortiOS server load balancing to load balance traffic to multiple backend servers.


You can configure FortiOS load balancing to intercept incoming traffic with a virtual server and distribute it among one or more backend real servers. By doing so, FortiOS enables multiple real servers to respond as if they were a single device or virtual server. This in turn means that more simultaneous requests can be handled by the servers.



Traffic can be balanced across multiple backend real servers based on a selection of load balancing methods including static (failover), round robin, weighted to account for different sized servers, or based on the health and performance of the server including round trip time, number of connections. The load balancer can balance layer 7 HTTP, HTTPS, SSL, generic layer 4 TCP, UDP and generic layer 3 IP protocols. Session persistence is supported based on injected HTTP/HTTPS cookies or the SSL session ID.

You can bind up to 8 real servers can to one virtual server. The real server topology is transparent to end users, and the users interact with the system as if it were only a single server with the IP address and port number of the virtual server. The real servers may be interconnected by high-speed LAN or by geographically dispersed WAN. The FortiGate unit schedules requests to the real servers and makes parallel services of the virtual server to appear to involve a single IP address.

There are additional benefits to load balancing. First, because the load is distributed across multiple servers, the service being provided can be highly available. If one of the servers breaks down, the load can still be handled by the other servers. Secondly, this increases scalability. If the load increases substantially, more servers can be added behind the FortiGate unit to cope with the increased load.

Server load balancing configuration

Traffic can be balanced across multiple backend real servers based on a selection of load balancing methods including static (failover), round robin, weighted to account for different sized servers, or based on the health and performance of the server including round trip time, number of connections. The load balancer can balance layer 7 HTTP, HTTPS, SSL, generic layer 4 TCP, UDP and generic layer 3 IP protocols. Session persistence is supported based on injected HTTP/HTTPS cookies or the SSL session ID.

You can bind up to 8 real servers can to one virtual server. The real server topology is transparent to end users, and the users interact with the system as if it were only a single server with the IP address and port number of the

virtual server. The real servers may be interconnected by high-speed LAN or by geographically dispersed WAN. The FortiGate unit schedules requests to the real servers and makes parallel services of the virtual server to appear to involve a single IP address.

Load balancing and other FortiOS features

Flow-based and proxy-based security features such as virus scanning, IPS, DLP, application control, and web filtering can be applied to load balanced sessions. This includes SSL offloading and multiplexing. Applying these UTM features to load balancing traffic may reduce load balancing performance.

Authentication is not supported for load balancing sessions. Usually FortiGate load balancing is used to allow public access to services on servers protected by a FortiGate unit. Authentication is not generally not required for this kind of configuration.

Features such web proxying, web caching, and WAN optimization also do not work with load balanced sessions. However, most other features that can be applied by a security policy are supported.

Configuring load balancing from the GUI

A virtual server is a specialized firewall virtual IP that performs server load balancing. From the GUI you add load balancing virtual server by going to **Policy & Objects > Virtual Servers**.

You can use the GUI to configure IPv, IPv6, IPv4 to IPv6 (NAT46), or IPv6 to IPv4 (NAT64) load balancing.

Type

Select the type of virtual server to configure. You can select **IPv4**, **IPv6**, **NAT46**, or **NAT64**. If Type is set to NAT46 or NAT64 you have fewer load balancing options (just HTTP, TCP, UDP and IP) and you can't configure advanced SSL and HTTPS load balancing features.

Name

Enter the name for the virtual server.

Type

Select the protocol to be load balanced by the virtual server. If you select a general protocol such as **IP**, **TCP**, or **UDP** the virtual server load balances all IP, TCP, or UDP sessions. If you select specific protocols such as **HTTP**, **HTTPS**, or **SSL** you can apply additional server load balancing features such as **Persistence** and **HTTP Multiplexing**.

- Select **HTTP** to load balance only HTTP sessions with destination port number that matches the **Virtual Server Port** setting. Change **Virtual Server Port** to match the destination port of the sessions to be load balanced (usually port 80 for HTTP sessions). You can also select **HTTP Multiplex**. You can also set **Persistence** to **HTTP Cookie** to select cookie-based persistence.
- Select **HTTPS** to load balance only HTTPS sessions with destination port number that matches the **Virtual Server Port** setting. Change **Virtual Server Port** to match the destination port of the sessions to be load balanced (usually port 443 for HTTPS sessions). You can also select **Multiplex HTTP requests/responses**. You can also set **Persistence** to **HTTP Cookie** to select cookie-based persistence. You can also set **Persistence** to **SSL Session ID**.

- Select **IMAPS** to load balance only IMAPS sessions with destination port number that matches the **Virtual Server Port** setting. Change **Virtual Server Port** to match the destination port of the sessions to be load balanced (usually port 993 for IMAPS sessions). You can also set **Persistence** to **SSL Session ID**.
- Select **POP3S** to load balance only POP3S sessions with destination port number that matches the **Virtual Server Port** setting. Change **Virtual Server Port** to match the destination port of the sessions to be load balanced (usually port 995 for POP3S sessions). You can also set **Persistence** to **SSL Session ID**.
- Select **SMTPS** to load balance only SMTPS sessions with destination port number that matches the **Virtual Server Port** setting. Change **Virtual Server Port** to match the destination port of the sessions to be load balanced (usually port 465 for SMTPS sessions). You can also set **Persistence** to **SSL Session ID**.
- Select **SSL** to load balance only SSL sessions with destination port number that matches the **Virtual Server Port** setting. Change **Virtual Server Port** to match the destination port of the sessions to be load balanced.
- Select **TCP** to load balance only TCP sessions with destination port number that matches the **Virtual Server Port** setting. Change **Virtual Server Port** to match the destination port of the sessions to be load balanced.
- Select **UDP** to load balance only UDP sessions with destination port number that matches the **Virtual Server Port** setting. Change **Virtual Server Port** to match the destination port of the sessions to be load balanced.
- Select **IP** to load balance all sessions accepted by the security policy that contains this virtual server.

Interface

Select the virtual server external or outgoing interface from the list. The outgoing interface is connected to the source network and receives the packets to be forwarded to the destination network.

Virtual Server IP

The IPv4 address of the virtual server. This is an IP address on the external interface that you want to map to an address on the destination network.

Virtual Server Port

Enter the external port number that you want to map to a port number on the destination network. Sessions with this destination port are load balanced by this virtual server.

Load Balance Method

Select the load balancing method used by the virtual server.

Persistence

Configure persistence to make sure that a user is connected to the same server every time they make a request that is part of the same session. Session persistence is supported for HTTP and SSL sessions.

Health Check

Select which health check monitor configuration will be used to determine a server's connectivity status.

HTTP Multiplexing

Select to use the FortiGate unit to multiplex multiple client connections into a few connections between the FortiGate unit and the real server.

Preserve Client IP

Select to preserve the IP address of the client in the `X-Forwarded-For` HTTP header. This can be useful if you want log messages on the real servers to the client's original IP address. If this option is not selected, the header will contain the IP address of the FortiGate unit.

This option appears only if **Type** is set to HTTP or HTTPS.

SSL Offloading

Accelerate clients' SSL connections to the server by using the FortiGate to perform SSL operations. This option appears only if **Type** is set to one of the SSL protocols.

Mode

Select which segments of the SSL connection will receive SSL offloading. You can select **Client <-> FortiGate** (or half mode) or **Full** (full mode).

This option appears only if **Type** is set to one of the SSL protocols.

Certificate

Select the certificate to use with **SSL Offloading**. The certificate key size must be 1024 or 2048 bits. 4096-bit keys are not supported.

This option appears only if **Type** is set to one of the SSL protocols.

Real Servers

Add Real Servers to the virtual server. The virtual server load balances traffic to these real servers. See [Real servers on page 2607](#).

Configuring load balancing from the CLI

From the CLI you configure IPv4 load balancing by adding a firewall virtual IP and setting the virtual IP type to server load balance:

```
config firewall vip
    edit Vserver-HTTP-1
        set type server-load-balance
    ...
```

Sever load balancing is also supported for:

- IPv6 using the command `config firewall vip6`
- IPv6 to IPv4 using the command `config firewall vip64`
- IPv4 to IPv6 using the commmand `config firewall vip46`

Configuration is the same as IPv4 VIPs except support for advanced HTTP and SSL related features is not available. IPv6 server load balancing supports all the same server types as IPv4 server load balancing (HTTP, HTTPS, IMAPS, POP3S, SMTPS, SSL, TCP, UDP, and IP). IPv4 to IPv6 and IPv6 to IPv4 server load balancing supports fewer server types (HTTP, TCP, UDP, and IP).

A virtual server includes a virtual server IP address bound to an interface. The virtual server IP address is the destination address incoming packets to be load balanced and the virtual server is bound to the interface that receives the packets to be load balanced.

For example, if you want to load balance incoming HTTP traffic from the Internet to a group of web servers on a DMZ network, the virtual server IP address is the known Internet IP address of the web servers and the virtual server binds this IP address to the FortiGate interface connected to the Internet.

When you bind the virtual server's external IP address to a FortiGate unit interface, by default, the network interface responds to ARP requests for the bound IP address. Virtual servers use proxy ARP, as defined in [RFC 1027](#), so that the FortiGate unit can respond to ARP requests on a network for a real server that is actually installed on another network. In some cases you may not want the network interface sending ARP replies. You can use the `arp-reply` option to disable sending ARP replies:

```
config firewall vip
  edit Vserver-HTTP-1
    set type server-load-balance
    set arp-reply disable
  ...
```

The load balancing virtual server configuration also includes the virtual server port. This is the TCP port on the bound interface that the virtual server listens for traffic to be load balanced on. The virtual server can listen on any port.

Load balancing methods

The load balancing method defines how sessions are load balanced to real servers. A number of load balancing methods are available as listed below.

All load balancing methods will not send traffic to real servers that are down or not responding. However, the FortiGate unit can only determine if a real server is not responding by using a health check monitor. You should always add at least one health check monitor to a virtual server or to individual real servers, or load balancing methods may attempt to distribute sessions to real servers that are not functioning.

Static

The traffic load is statically spread evenly across all real servers. However, sessions are not assigned according to how busy individual real servers are. This load balancing method provides some persistence because all sessions from the same source address always go to the same real server. However, the distribution is stateless, so if a real server is added or removed (or goes up or down) the distribution is changed and persistence could be lost.

Round Robin

Directs new requests to the next real server, and treats all real servers as equals regardless of response time or number of connections. Dead real servers or non responsive real servers are avoided.

Weighted

Real servers with a higher weight value receive a larger percentage of connections. Set the real server weight when adding a real server.

Least Session

Directs requests to the real server that has the least number of current connections. This method works best in environments where the real servers or other equipment you are load balancing all have similar capabilities. This load balancing method uses the FortiGate session table to track the number of sessions being processed by each real server. The FortiGate unit cannot detect the number of sessions actually being processed by a real server.

Least RTT

Directs sessions to the real server with the least round trip time. The round trip time is determined by a Ping health check monitor and is defaulted to 0 if no Ping health check monitors are added to the virtual server.

First Alive

Always directs sessions to the first alive real server. This load balancing schedule provides real server failover protection by sending all sessions to the first alive real server and if that real server fails, sending all sessions to the next alive real server. Sessions are not distributed to all real servers so all sessions are processed by the “first” real server only.

First refers to the order of the real servers in the virtual server configuration. For example, if you add real servers A, B and C in that order, then all sessions always go to A as long as it is alive. If A goes down then sessions go to B and if B goes down sessions go to C. If A comes back up sessions go back to A. Real servers are ordered in the virtual server configuration in the order in which you add them, with the most recently added real server last. If you want to change the order you must delete and re-add real servers in the required order.

HTTP Host

Load balances HTTP host connections across multiple real servers using the host's HTTP header to guide the connection to the correct real server.

Session persistence

Use persistence to make sure that a user is connected to the same real server every time they make an HTTP, HTTPS, or SSL request that is part of the same user session. For example, if you are load balancing HTTP and HTTPS sessions to a collection of eCommerce web servers, when a user is making a purchase they will be starting multiple sessions as they navigate the eCommerce site. In most cases all of the sessions started by this user during on eCommerce session should be processed by the same real server. Typically, the HTTP protocol keeps track of these related sessions using cookies. HTTP cookie persistence makes sure that all sessions that are part of the same user session are processed by the same real server

When you configure persistence, the FortiGate unit load balances a new session to a real server according to the load balance method. If the session has an HTTP cookie or an SSL session ID, the FortiGate unit sends all subsequent sessions with the same HTTP cookie or SSL session ID to the same real server. For more information about HTTP and HTTPS persistence, see [“HTTP and HTTPS persistence”](#).

Real servers

Add real servers to a load balancing virtual server to provide the information the virtual server requires to be able to send sessions to the server. A real server configuration includes the IP address of the real server and port number that the real server receives sessions on. The FortiGate unit sends sessions to the real server's IP address using the destination port number in the real server configuration.

When configuring a real server you can also specify the weight (used if the load balance method is set to weighted) and you can limit the maximum number of open connections between the FortiGate unit and the real server. If the maximum number of connections is reached for the real server, the FortiGate unit will automatically switch all further connection requests other real servers until the connection number drops below the specified limit. Setting Maximum Connections to 0 means that the FortiGate unit does not limit the number of connections to the real server.

Real server active, standby, and disabled modes

By default the real server mode setting is active indicating that the real server is available to receive connections. If the real server is removed from the network (for example, for routine maintenance or because of a hardware or software failure) you can change the mode to standby or disabled. In disabled mode the FortiGate unit no longer sends sessions to the real server.

If a real server is in standby mode the FortiGate also does not send sessions to it unless other real servers added to the same virtual server become unavailable. For example:

- A virtual server that includes two real servers one in active mode and one in standby mode. If the real server in active mode fails, the real server in standby mode is changed to active mode and all sessions are sent to this real server.
- A virtual server includes three real servers, two in active mode and one in standby mode, if one of the real servers in active mode fails, the real server in standby mode is changed to active mode and sessions are load balanced between it and still operating real server. If both real servers in active mode fail, all sessions are sent to the real server in standby mode.

Adding real servers from the GUI

To add a real server from the GUI go to **Policy & Objects > Virtual Servers**, edit a virtual server and under **Real Servers** select **Create New** to add a real server to this virtual server.

IP Address

Enter the IP address of the real server.

Port

Enter the port number on the destination network to which the external port number is mapped.

Weight

Enter the weight value of the real server. The higher the weight value, the higher the percentage of connections the server will handle. A range of 1-255 can be used. This option is available only if the associated virtual server's load balance method is **Weighted**.

Max Connections

Enter the limit on the number of active connections directed to a real server. A range of 1-99999 can be used. If the maximum number of connections is reached for the real server, the FortiGate unit will automatically switch all further connection requests to another server until the connection number drops below the specified limit.

Setting **Maximum Connections** to **0** means that the FortiGate unit does not limit the number of connections to the real server.

HTTP Host

Enter the HTTP header for load balancing across multiple real servers. This feature is used for load balancing HTTP host connections across multiple real servers using the host's HTTP header to guide the connection to the correct real server, providing better load balancing for those specific connections.

Mode

Select a mode for the real server. The real server can be active, on standby, or disabled.

Adding real servers from the CLI

To add a real server from the CLI you configure a virtual server and add real servers to it. For example, to add three real servers to a virtual server that load balances UDP sessions on port 8190 using weighted load balancing. For each real server the port is not changed. The default real server port is 0 resulting in the traffic being sent the real server with destination port 8190. Each real sever is given a different weight. Servers with higher weights have a max-connections limit to prevent too many sessions from being sent to them.

```
config firewall vip
  edit Vserver-UDP-1
    set type server-load-balance
    set server-type udp
    set ldb-method weighted
    set extip 172.20.120.30
    set extintf wan1
    set extport 8190
    set monitor ping-mon-1
    config realservers
      edit 1
        set ip 10.31.101.30
        set weight 100
        set max-connections 10000
      next
      edit 2
        set ip 10.31.101.40
        set weight 100
        set max-connections 10000
      next
      edit 3
        set ip 10.31.101.50
        set weight 10
      end
    end
end
```

Health check monitoring

From the FortiGate GUI you can go to **Policy & Objects > Health Check** and configure health check monitoring so that the FortiGate unit can verify that real servers are able respond to network connection attempts. If a real server responds to connection attempts the load balancer continues to send sessions to it. If a real server stops responding to connection attempts the load balancer assumes that the server is down and does not send sessions to it. The health check monitor configuration determines how the load balancer tests the real servers. You can use a single health check monitor for multiple load balancing configurations.

You can configure TCP, HTTP and Ping health check monitors. Usually you would want the health check monitor to use the same protocol for checking the health of the server as the traffic being load balanced to it. For example, for an HTTP load balancing configuration you would normally use an HTTP health check monitor.

For the TCP and HTTP health check monitors you can specify the destination port to use to connect to the real servers. If you set the port to 0, the health check monitor uses the port defined in the real server. This allows you to use the same health check monitor for multiple real servers using different ports. You can also configure the interval, timeout and retry. A health check occurs every number of seconds indicated by the interval. If a reply is

not received within the timeout period the health check is repeated every second. If no response is received after the number of configured retries, the virtual server is considered unresponsive, and load balancing does not send traffic to that real server. The health check monitor will continue to contact the real server and if successful, the load balancer can resume sending sessions to the recovered real server.

The default health check configuration has an interval of 10 seconds, a timeout of 2 seconds and a retry of 3. This means that the health check monitor checks the health of a real server every 10 seconds. If a reply is not received within 2 seconds the health check monitor re-checks the server every second for 3 retries. If no response is received for 2 seconds after the final retry the server is considered unresponsive. This entire process takes a total of 7 seconds to consider a virtual server as unresponsive (2 second timeout + (3 re-checks x 1 second) + 2 second timeout = 7 seconds). Since this health check process is repeated every 10 seconds, a server can be down for a maximum of $10 + 7 = 17$ seconds before the health check monitor considers it down.

For HTTP health check monitors, you can add URL that the FortiGate unit connects to when sending a get request to check the health of a HTTP server. The URL should match an actual URL for the real HTTP servers. The URL is optional.

The URL would not usually include an IP address or domain name. Instead it should start with a "/" and be followed by the address of an actual web page on the real server. For example, if the IP address of the real server is 10.31.101.30, the URL "/test_page.htm" causes the FortiGate unit to send an HTTP get request to "http://10.31.101.30/test_page.htm".

For HTTP health check monitors, you can also add a matched content phrase that a real HTTP server should include in response to the get request sent by the FortiGate unit using the content of the URL option. If the URL returns a web page, the matched content should exactly match some of the text on the web page. You can use the URL and Matched Content options to verify that an HTTP server is actually operating correctly by responding to get requests with expected web pages. Matched content is only required if you add a URL.

For example, you can set matched content to "server test page" if the real HTTP server page defined by the URL option contains the phrase "server test page". When the FortiGate unit receives the web page in response to the URL get request, the system searches the content of the web page for the matched content phrase.

Name

Enter the name of the health check monitor configuration.

Type

Select the protocol used to perform the health check.

- TCP
- HTTP
- PING

Port

Enter the port number used to perform the health check. If you set the **Port** to 0, the health check monitor uses the port defined in the real server. This way you can use a single health check monitor for different real servers.

This option does not appear if the **Type** is **PING**.

Interval

Enter the number of seconds between each server health check.

URL

For HTTP health check monitors, add a URL that the FortiGate unit uses when sending a get request to check the health of a HTTP server. The URL should match an actual URL for the real HTTP servers. The URL is optional.

The URL would not usually include an IP address or domain name. Instead it should start with a “/” and be followed by the address of an actual web page on the real server. For example, if the IP address of the real server is 10.10.10.1, the **URL** “/test_page.htm” causes the FortiGate unit to send an HTTP get request to “http://10.10.10.1/test_page.htm”.

This option appears only if **Type** is **HTTP**.

Matched Content

For HTTP health check monitors, add a phrase that a real HTTP server should include in response to the get request sent by the FortiGate unit using the content of the **URL** option. If the **URL** returns a web page, the **Matched Content** should exactly match some of the text on the web page. You can use the **URL** and **Matched Content** options to verify that an HTTP server is actually operating correctly by responding to get requests with expected web pages. Matched content is only required if you add a URL.

For example, you can set **Matched Content** to “server test page” if the real HTTP server page defined by the URL option contains the phrase “server test page”. When the FortiGate unit receives the web page in response to the URL get request, the system searches the content of the web page for the **Matched Content** phrase.

This option appears only if **Type** is **HTTP**.

Max Redirects

For an HTTP health check monitor, specify the maximum number of redirects that the health check monitor will follow when testing the health of the real HTTP server. This feature allows you to do health checking of the HTTP server is accessed through one or more redirects.

Timeout

Enter the number of seconds which must pass after the server health check to indicate a failed health check.

Retry

Enter the number of times, if any, a failed health check will be retried before the server is determined to be inaccessible.

Load balancing limitations

The following limitations apply when adding virtual IPs, load balancing virtual servers, and load balancing real servers. Load balancing virtual servers are actually server load balancing virtual IPs. You can add server load balance virtual IPs from the CLI.

- Virtual IP **External IP Address/Range** entries or ranges cannot overlap with each other or with load balancing virtual server **Virtual Server IP** entries.
- A virtual IP **Mapped IP Address/Range** cannot be 0.0.0.0 or 255.255.255.255.
- A real server **IP** cannot be 0.0.0.0 or 255.255.255.255.
- If a static NAT virtual IP **External IP Address/Range** is 0.0.0.0, the **Mapped IP Address/Range** must be a single IP address.

- If a load balance virtual IP **External IP Address/Range** is 0.0.0.0, the **Mapped IP Address/Range** can be an address range.
- When port forwarding, the count of mapped port numbers and external port numbers must be the same. The GUI does this automatically but the CLI does not.
- Virtual IP and virtual server names must be different from firewall address or address group names.

Monitoring load balancing

From the GUI you can go to **Monitor > Load Balance Monitor** to monitor the status of configured virtual servers and real servers and start or stop the real servers. You can also use the `get test ipldb` command from the CLI to display similar information.

For each real server the monitor displays health status (up or down), active sessions, round trip time (RTT) and the amount of bytes of data processed. From the monitor page you can also stop sending new sessions to any real server. When you select to stop sending sessions the FortiGate unit performs a graceful stop by continuing to send data for sessions that were established or persistent before you selected stop. However, no new sessions are started.

Real Server

The IP addresses of the existing real servers.

Status

Displays the health status according to the health check results for each real server. A green arrow means the server is up. A red arrow means the server is down.

Mode

The mode of the health check monitor. Can be active, standby, or disabled.

Monitor Events

Display each real server's up and down times.

Active Sessions

Display each real server's active sessions.

RTT (ms)

Displays the Round Trip Time (RTT) of each real server. By default, the RTT is "<1". This value will change only when ping monitoring is enabled on a real server.

Bytes Processed

Displays the traffic processed by each real server.

Graceful Stop/Start

Select to start or stop real servers. When stopping a server, the FortiGate unit will not accept new sessions but will wait for the active sessions to finish.

Load balancing diagnose commands

You can also use the following diagnose commands to view status information for load balancing virtual servers and real servers:

```
diagnose firewall vip realserver {down | healthcheck | list | up}
diagnose firewall vip virtual-server {filter | real-server | stats}
```

For example, the following command lists and displays status information for all real servers:

```
diagnose firewall vip virtual-server real-server

vd root/0 vs vs/2 addr 10.31.101.30:80 status 1/1
conn: max 0 active 0 attempts 0 success 0 drop 0 fail 0

vd root/0 vs vs/2 addr 10.31.101.20:80 status 1/1
conn: max 0 active 0 attempts 0 success 0 drop 0 fail 0
```

Many of the diagnostic commands involve retrieving information about one or more virtual servers. To control which servers are queried you can define a filter:

```
diagnose firewall vip virtual-server filter <filter_str>
```

Where <filter_str> can be:

- **clear** erase the current filter
- **dst** the destination address range to filter by
- **dst-port** the destination port range to filter by
- **list** display the current filter
- **name** the vip name to filter by
- **negate** negate the specified filter parameter
- **src** the source address range to filter by
- **src-port** the source port range to filter by
- **vd** index of virtual domain. -1 matches all

The default filter is empty so no filtering is done.

Logging diagnostics

The logging diagnostics provide information about two separate features:

```
diagnose firewall vip virtual-server filter
```

- **filter** sets a filter for the virtual server debug log
- The filter option controls what entries the virtual server daemon will log to the console if **diagnose debug application vs** level is non-zero. The filtering can be done on source, destination, virtual-server name, virtual domain, and so on:

```
diagnose firewall vip virtual-server filter <filter_str>
```

Where <filter_str> can be

- **clear** erase the current filter
- **dst** the destination address range to filter by
- **dst-port** the destination port range to filter by
- **list** display the current filter
- **name** the virtual-server name to filter by
- **negate** negate the specified filter parameter

- `src` the source address range to filter by
- `src-port` the source port range to filter by
- `vd` index of virtual domain. -1 matches all

The default filter is empty so no filtering is done.

Real server diagnostics

Enter the following command to list all the real servers:

```
diagnose firewall vip virtual-server real-server list
```

In the following example there is only one virtual server called `slb` and it has two real-servers:

```
diagnose firewall vip virtual-server server
vd root/0 vs slb/2 addr 172.16.67.191:80 status 1/1
conn: max 10 active 0 attempts 0 success 0 drop 0 fail 0
http: available 0 total 0

vd root/0 vs slb/2 addr 172.16.67.192:80 status 1/1
conn: max 10 active 1 attempts 4 success 4 drop 0 fail 0
http: available 1 total 1
```

The `status` indicates the administrative and operational status of the real-server.

- `max` indicates that the real-server will only allow 10 concurrent connections.
- `active` is the number of current connections to the server attempts is the total number of connections attempted success is the total number of connections that were successful.
- `drop` is the total number of connections that were dropped because the active count hit max.
- `fail` is the total number of connections that failed to complete due to some internal problem (for example, lack of memory).

If the virtual server has HTTP multiplexing enabled then the HTTP section indicates how many established connections to the real-server are available to service a HTTP request and also the total number of connections.

HTTP and HTTPS load balancing, multiplexing, and persistence

In a firewall load balancing virtual server configuration, you can select HTTP to load balance only HTTP sessions. The virtual server will load balance HTTP sessions received at the virtual server interface with destination IP address that matches the configured virtual server IP and destination port number that matches the configured virtual server port. The default virtual server port for HTTP load balancing is 80, but you can change this to any port number. Similarly for HTTPS load balancing, set the virtual server type to HTTPS and then select the interface, virtual server IP, and virtual server port that matches the HTTPS traffic to be load balanced. Usually HTTPS traffic uses port 443.

You can also configure load balancing to offload SSL processing for HTTPS and SSL traffic. See [SSL/TLS load balancing on page 2619](#).

HTTP and HTTPS multiplexing

For both HTTP and HTTPS load balancing you can multiplex HTTP requests and responses over a single TCP connection. HTTP multiplexing is a performance saving feature of HTTP/1.1 compliant web servers that provides the ability to pipeline many unrelated HTTP or HTTPS requests on the same connection. This allows a single HTTPD process on the server to interleave and serve multiple requests. The result is fewer idle sessions on the

web server so server resources are used more efficiently. HTTP multiplexing can take multiple separate inbound sessions and multiplex them over the same internal session. This may reduce the load on the backend server and increase the overall performance.

HTTP multiplexing may improve performance in some cases. For example, if users web browsers are only compatible with HTTP 1.0. HTTP multiplexing can also improve performance between a web server and the FortiGate unit if the FortiGate unit is performing SSL acceleration. However, in most cases HTTP multiplexing should only be used if enabling it leads to a measurable improvement in performance.

To enable HTTP multiplexing from the GUI, select multiplex HTTP requests/responses over a single TCP connection. To enable HTTP multiplexing from the CLI enable the `http-multiplex` option.

Preserving the client IP address

Select preserve client IP from the GUI or enable the `http-ip-header` option from the CLI to preserve the IP address of the client in the `X-Forwarded-For` HTTP header. This can be useful in an HTTP multiplexing configuration if you want to be able to see the original client IP address in log messages on the destination web server. If this option is not selected, the `X-Forwarded-For` HTTP header contains the IP address of the FortiGate unit.

Preserving the client IP address in a different HTTP header

If you select preserve client IP from the GUI or enable the `http-ip-header` option from the CLI you can also preserve the client IP in a different HTTP header. This can be useful if you want to use a custom header name instead of `X-Forwarded-For`.

You can add the custom header name from the CLI. When `http-ip-header` is enabled you can add a custom header name to the `http-ip-header-name` option. If you don't add a name the `X-Forwarded-For` header is used.

HTTP and HTTPS persistence

Configure load balancing persistence for HTTP or HTTPS to make sure that a user is connected to the same server every time they make a request that is part of the same session. HTTP cookie persistence uses injected cookies to enable persistence.

When you configure persistence, the FortiGate unit load balances a new session to a real server according to the **Load Balance Method**. If the session has an HTTP cookie or an SSL session ID, the FortiGate unit sends all subsequent sessions with the same HTTP cookie or SSL session ID to the same real server.

The following example shows how to enable cookie persistence and set the cookie domain to `.example.org`.

```
config firewall vip
  edit HTTP_Load_Balance
    set type server-load-balance
    set server-type http
    set extport 8080
    set extintf port2
    set extip 192.168.20.20
    set persistence http-cookie
    set http-cookie-domain .example.org
    config realservers
    edit 1
      set ip 10.10.10.1
      set port 80
    next
```

```

edit 2
    set ip 10.10.10.2
    set port 80
next
edit 3
    set ip 10.10.10.3
    set port 80
end

```

How HTTP cookie persistence options work

The following options are available for the `config firewall vip` command when `type` is set to `server-load-balance`, `server-type` is set to `http` or `https` and `persistence` is set to `http-cookie`:

```

http-cookie-domain-from-host
http-cookie-domain
http-cookie-path
http-cookie-generation
http-cookie-age
http-cookie-share
https-cookie-share

```

When HTTP cookie persistence is enabled the FortiGate unit inserts a header of the following form into each HTTP response unless the corresponding HTTP request already contains a `FGTServer` cookie:

```
Set-Cookie: FGTServer=E7D01637C4B08E89A6714213A9D85D9C7E4D8158; Version=1; Max-Age=3600
```

The value of the `FGTServer` cookie encodes the server that traffic should be directed to. The value is encoded so as to not leak information about the internal network.

Enable `http-cookie-domain-from-host` to extract the cookie domain from the `host:` header in the HTTP request. For example, to restrict the cookie to `.server.com`, enter:

The generated cookies could have the following form if the **Host:** header contains **exhost.com**:

```
Set-Cookie: FGTServer=E7D01637C4B08E89A6714213A9D85D9C7E4D8158; Version=1;
Domain=.exhost.com; Max-Age=3600
```

For more information, see [“HTTP host-based load balancing”](#).

Use `http-cookie-domain` to restrict the domain that the cookie should apply to. For example, to restrict the cookie to `.server.com`, enter:

```
set http-cookie-domain .server.com
```

Now all generated cookies will have the following form:

```
Set-Cookie: FGTServer=E7D01637C4B08E89A6714213A9D85D9C7E4D8158; Version=1;
Domain=.server.com; Max-Age=3600
```

Use `http-cookie-path` to limit the cookies to a particular path. For example, to limit cookies to the path `/sales`, enter:

```
set http-cookie-path /sales
```

Now all generated cookies will have the following form:

```
Set-Cookie: FGTServer=E7D01637C4B08E89A6714213A9D85D9C7E4D8158; Version=1;
Domain=.server.com; Path=/sales; Max-Age=3600
```

Use `http-cookie-age` to change how long the browser caches the cookie. You can enter an age in minutes or set the age to 0 to make the browser keep the cookie indefinitely:

```
set http-cookie-age 0
```

Now all generated cookies will have the following form:

```
Set-Cookie: FGTSer=E7D01637C4B08E89A6714213A9D85D9C7E4D8158; Version=1;
Domain=.server.com; Path=/sales
```

Use `http-cookie-generation` to invalidate all cookies that have already been generated. The exact value of the generation is not important, only that it is different from any generation that has already been used for cookies in this domain. The simplest approach is to increment the generation by one each time invalidation is required. Since the default is 0, enter the following to invalidate all existing cookies:

```
set http-cookie-generation 1
```

Use `http-cookie-share {disable | same-ip}` to control the sharing of cookies across virtual servers in the same virtual domain. The default setting `same-ip` means that any `FGTSer` cookie generated by one virtual server can be used by another virtual server in the same virtual domain. For example, if you have an application that starts on HTTP and then changes to HTTPS and you want to make sure that the same server is used for the HTTP and HTTPS traffic then you can create two virtual servers, one for port 80 (for HTTP) and one for port 443 (for HTTPS). As long as you add the same real servers to both of these virtual servers (and as long as both virtual servers have the same number of real servers with the same IP addresses), then cookies generated by accessing the HTTP server are reused when the application changes to the HTTPS server.

If for any reason you do not want this sharing to occur then select `disable` to make sure that a cookie generated for a virtual server cannot be used by other virtual servers.

Use `https-cookie-secure` to enable or disable using secure cookies. Secure cookies are disabled by default because secure cookies can interfere with cookie sharing across HTTP and HTTPS virtual servers. If enabled, then the `Secure` tag is added to the cookie inserted by the FortiGate unit:

```
Set-Cookie: FGTSer=E7D01637C4B08E89A6714213A9D85D9C7E4D8158; Version=1; Max-Age=3600;
Secure
```

HTTP host-based load balancing

When configuring HTTP or HTTPS load balancing you can select HTTP host load balancing to load balance HTTP host connections across multiple real servers using the host's HTTP header to guide the connection to the correct real server. HTTP 1.1 includes the concept of a virtual server which allows a HTTP or HTTPS server with a single external IP address to serve requests for multiple DNS domains by using the mandatory `Host:` header in a HTTP request to indicate which DNS domain the request is destined for.

FortiOS can load-balance HTTP and HTTPS connections among multiple real servers using the `Host:` header to guide the connection to the correct real server. The host load balancing method allows a real server to specify a `http-host` attribute which is the domain name of the traffic for that real server. Each real server can only specify a single domain name. The same domain name can appear in more than one real server but only the first one that is up will be used, any others are purely for redundancy. If the `Host:` header contains a domain that does not match any `http-host` entry then the connection will be dropped. A real server with no `http-host` can be matched by any `Host:` domain.

For example, consider a FortiGate unit that is load-balancing traffic to three real servers. Traffic for `www.example1.com` should go to `192.168.2.1`, traffic for `www.example2.com` should go to `192.168.2.2` and traffic to any other domain should go to `192.168.2.3`. To enable this configuration you would add a virtual server and set the load balance method to HTTP host. Then you would add three real servers and set the HTTP host of the real server with IP address `192.168.2.1` to `www.example1.com`, the HTTP host of the real server with IP address `192.168.2.2` to `www.example2.com` and you would not specify an HTTP host for the third real server.

The configuration of a virtual IP to achieve this result would be:

```
config firewall vip
edit "http-host-ldb"
set type server-load-balance
```



```
set extip 172.16.67.195
set extintf "lan"
set server-type http
set ldb-method http-host
set extport 80
config realservers
edit 1
    set http-host "www.example1.com"
    set ip 192.168.2.1
    set port 80
next
edit 2
    set http-host "www.example2.com"
    set ip 192.168.2.2
    set port 80
next
edit 3
    set ip 192.168.2.3
    set port 80
next
end
```

Host load balancing and HTTP cookie persistence

In an HTTP host-based load balancing configuration with HTTP cookie persistence enabled you can optionally configure cookie persistence to use the domain set in the host header as the cookie domain. You can do this by enabling the `http-cookie-domain-from-host` option, for example:

```
config firewall vip
edit "http-host-ldb"
    set type server-load-balance
    set extip 172.16.67.195
    set extintf "lan"
    set server-type http
    set ldb-method http-host
    set extport 80
    set persistence http-cookie
    set http-cookie-domain-from-host enable
    config realservers
        edit 1
            set http-host "www.example1.com"
            set ip 192.168.2.1
            set port 80
        next
        edit 2
            set http-host "www.example2.com"
            set ip 192.168.2.2
            set port 80
        next
        edit 3
            set ip 192.168.2.3
            set port 80
        next
    end
end
```

SSL/TLS load balancing

In a firewall load balancing virtual server configuration, you can select SSL to load balance only SSL and TLS sessions. The virtual server will load balance SSL and TLS sessions received at the virtual server interface with destination IP address that matches the configured virtual server IP and destination port number that matches the configured virtual server port. Change this port to match the destination port of the sessions to be load balanced.

For SSL load balancing you can also set persistence to SSL session ID. Persistence is achieved by the FortiGate unit sending all sessions with the same SSL session ID to the same real server. When you configure persistence, the FortiGate unit load balances a new session to a real server according to the **Load Balance Method**. If the session has an SSL session ID, the FortiGate unit sends all subsequent sessions with the same SSL session ID to the same real server.


SSL/TLS offloading

Use SSL offloading to accelerate clients' SSL or HTTPS connections to real servers by using the FortiGate unit to perform SSL/TLS operations (offloading them from the real servers using the FortiGate unit's SSL acceleration hardware). FortiGate units can offload most versions of SSL/TLS, including SSL 3.0, TLS 1.0 and TLS 1.2. SSL/TLS offloading is available on FortiGate units that support SSL acceleration.

To configure SSL offloading from the GUI go to **Policy & Objects > Virtual Servers**. Add a virtual server and set the type to HTTPS or SSL and select the SSL offloading type (**Client <-> FortiGate** or **Full**).

Select **Client <-> FortiGate** to apply hardware accelerated SSL/TLS processing only to the part of the connection between the client and the FortiGate unit. This mode is called half mode SSL offloading. The segment between the FortiGate unit and the server will use clear text communications. This results in best performance, but cannot be used in failover configurations where the failover path does not have an SSL accelerator.

Select **Full** to apply hardware accelerated SSL processing to both parts of the connection: the segment between client and the FortiGate unit, and the segment between the FortiGate unit and the server. The segment between the FortiGate unit and the server uses encrypted communications, but the handshakes are abbreviated. This is not as efficient as half mode SSL offloading, but still improves performance. As well, full-mode SSL offloading can be used in failover configurations where the failover path does not have an SSL accelerator. If the server is already configured to use SSL, this also enables SSL acceleration without requiring changes to the server's configuration.

SSL Offloading modes (Half Mode and Full Mode)

Configuring SSL offloading also requires selecting a certificate to use for the SSL offloading sessions. SSL offloading supports key sizes up to 4096. FortiGate models with CP9 processors support 3072 and 4096 DH bit sizes in hardware. All FortiGate models up to and including those with CP8 processors only support offloading DH bit sizes up to 2048 so any sizes larger than that are done in software and thus are relatively resource intensive.

The following CLI command shows an example half mode HTTPS SSL offloading configuration. In the example the `ssl-mode` option sets the SSL offload mode to `half` (which is the default mode).

```
config firewall vip
  edit Vserver-ssl-offload
    set type server-load-balance
    set server-type https
    set ldb-method round-robin
    set extip 172.20.120.30
    set extintf wan1
    set extport 443
```

```
set persistence ssl-session-id
set ssl-mode half
set ssl-certificate my-cert
set monitor tcp-mon-1
config realservers
  edit 1
    set ip 10.31.101.30
    set port 443
  next
  edit 2
    set ip 10.31.101.40
    set port 443
  end
end
```

Separate virtual-server client and server TLS version and cipher configuration

In some cases, you may want the to use different versions of SSL or TLS on the client to FortiGate connection than on the FortiGate to server connection. For example, you may want to use the FortiGate to protect a legacy SSL 3.0 or TLS 1.0 server while making sure that client to FortiGate connections must always use the higher level of protection offered by TLS 1.1 or greater. Also, in some cases you might want to protect a server that only has weak ciphers (for example, DES or RC4) while making sure that all connections between the FortiGate and the client use a strong cipher for better protection.

The following options are available when configuring server load balancing for HTTPS sessions configured with the following command:

```
config firewall vip
  edit server-name
    set type server-load-balance
    set server-type https
    set ssl-mode full
  ...
```

Setting the SSL/TLS versions to use for server and client connections

The `ssl-server-min-version`, `ssl-server-max-version`, `ssl-min-version` and `ssl-max-version` configuration options allow the minimum and maximum SSL/TLS versions for the client to FortiGate connection to be independent of the FortiGate to server configuration. By default these options are both set to `client` and the configured `ssl-min-version` and `ssl-max-version` settings are applied to both the client and the server connection.

You can change the `ssl-server-min-version` and `ssl-server-max-version` to apply different options to the server connection. The `ssl-min-version` and `ssl-max-version` settings are still applied to the client connection. If you set the `ssl-server-min-version` and `ssl-server-max-version` to an explicit version then both must be set to an explicit version.

The `ssl-server-min-version` and `ssl-server-max-version` options allow you to specify the minimum and maximum SSL/TLS versions the FortiGate will offer to the server (in the record header of the ClientHello) when performing full mode SSL offloading and thus the minimum and maximum SSL/TLS versions the FortiGate accepts from the server (in a ServerHello). If the server responds with a version in its ServerHello that is lower than `ssl-server-min-version` or higher than the `ssl-server-max-version` then the FortiGate terminates the connection.

Command syntax is:

```

config firewall vip
edit server-name
set type server-load-balance
set server-type https
set ssl-mode full
set ssl-min-version {ssl-3.0 | tls-1.0 | tls-1.1 | tls-1.2}
set ssl-max-version {ssl-3.0 | tls-1.0 | tls-1.1 | tls-1.2}
set ssl-server-min-version {ssl-3.0 | tls-1.0 | tls-1.1 | tls-1.2 | client}
set ssl-server-max-version {ssl-3.0 | tls-1.0 | tls-1.1 | tls-1.2 | client}

```

Setting the SSL/TLS cipher choices for server and client connections

The `ssl-algorithm` and `ssl-server-algorithm` configuration options allow the cipher choice for the FortiGate to server connection to be independent of the client to FortiGate connection. By default, `ssl-server-algorithm` is set to `client` and the configured `ssl-algorithm` setting is applied to both the client and the server connection.

You can change the `ssl-server-algorithm` to apply different options to the server connection. The `ssl-algorithm` setting is still applied to the client connection.

The following `ssl-server-algorithm` options are available:

- `high`, offer AES or 3DES cypher suites in the ServerHello
- `medium`, use AES, 3DES, or RC4 cypher suites in the ServerHello
- `low`, use AES, 3DES, RC4, or DES cypher suites in the ServerHello
- `custom`, specify custom cypher suites using the `config ssl-server-cipher-suites` and offer these custom cypher suites in the ServerHello.
- `client`, offer the cypher suites in the ServerHello that are offered in the ClientHello.

Command syntax is:

```

config firewall vip
edit server-name
set type server-load-balance
set server-type https
set ssl-mode full
set ssl-algorithm {high | medium | low | custom}
set ssl-server-algorithm {high | medium | low | custom | client}

```

If you set `ssl-server-algorithm` to `custom`, the syntax is:

```

config firewall vip
edit server-name
set type server-load-balance
set server-type https
set ssl-mode full
set ssl-server-algorithm custom
config ssl-server-cipher-suites
edit 10
set cipher <cipher-suite>
set versions {ssl-3.0 | tls-1.0 | tls-1.1 | tls-1.2}
next
edit 20
set cipher <cipher-suite>
set versions {ssl-3.0 | tls-1.0 | tls-1.1 | tls-1.2}
end

```

Protection from TLS protocol downgrade attacks

The `ssl-client-fallback` option, when enabled (the default configuration), performs downgrade attack prevention (RFC 7507).

Command syntax is:

```
config firewall vip
edit server-name
set type server-load-balance
set server-type https
set ssl-client-fallback {disable | enable}
```

Setting 3072- and 4096-bit Diffie-Hellman values

The `ssl-dh-bits` option allows you to specify the number of bits of the prime number used in the Diffie-Hellman exchange for RSA encryption of the SSL connection. Larger prime numbers are associated with greater cryptographic strength. You can set DH values from 768 to 4096 bits.

Command syntax is:

```
config firewall vip
edit server-name
set type server-load-balance
set server-type https
set ssl-dh-bits {768 | 1024 | 1536 | 2048 | 3072 | 4096}
```

Setting the DH bits to 2048 only provides the equivalent of a symmetric cipher in the range of 112 - 128 bits. This means that if AES 256 is used then the weakest point is the DH of 2048 and a value of at least 3072 should be used if the goal is to have 256 bits of security.

FortiGate models with CP9 processors support 3072 and 4096 DH bit sizes in hardware. All FortiGate models up to and including those with CP8 processors only support offloading DH bit sizes up to 2048 so any sizes larger than that are done in software and thus are relatively resource intensive.

Additional SSL load balancing and SSL offloading options

The following SSL load balancing and SSL offloading options are only available from the CLI:

```
ssl-client-session-state-max <sessionstates_int>
```

Enter the maximum number of SSL session states to keep for the segment of the SSL connection between the client and the FortiGate unit.

```
ssl-client-session-state-timeout <timeout_int>
```

Enter the number of minutes to keep the SSL session states for the segment of the SSL connection between the client and the FortiGate unit.

```
ssl-client-session-state-type {both | client | disable | time}
```

Select which method the FortiGate unit should use when deciding to expire SSL sessions for the segment of the SSL connection between the client and the FortiGate unit.

- **both:** Select to expire SSL session states when either `ssl-client-session-state-max` or `ssl-client-session-state-timeout` is exceeded, regardless of which occurs first.
- **count:** Select to expire SSL session states when `ssl-client-session-state-max` is exceeded.

- **disable:** Select to keep no SSL session states.
- **time:** Select to expire SSL session states when `ssl-client-session-state-timeout` is exceeded.

```
ssl-http-location-conversion {enable | disable}
```

Select to replace `http` with `https` in the reply's `Location` HTTP header field. For example, in the reply, `Location: http://example.com/` would be converted to `Location: https://example.com/`

```
ssl-http-match-host {enable | disable}
```

Enable (the default) to apply `Location` conversion to the reply's HTTP header only if the host name portion of `Location` matches the request's `Host` field, or, if the `Host` field does not exist, the host name portion of the request's URI.

If disabled, conversion occurs regardless of whether the host names in the request and the reply match.

For example, if host matching is enabled, and a request contains `Host: example.com` and the reply contains `Location: http://example.cc/`, the `Location` field does not match the host of the original request and the reply's `Location` field remains unchanged. If the reply contains `Location: http://example.com/`, however, then the FortiGate unit detects the matching host name and converts the reply field to `Location: https://example.com/`.

This option appears only if `ssl-http-location-conversion` is `enable`.

```
ssl-send-empty-frags {enable | disable}
```

Select to precede the record with empty fragments to protect from attacks on CBC IV. You might disable this option if SSL acceleration will be used with an old or buggy SSL implementation which cannot properly handle empty fragments.

```
ssl-server-session-state-max <sessionstates_int>
```

Enter the maximum number of SSL session states to keep for the segment of the SSL connection between the server and the FortiGate unit.

```
ssl-server-session-state-timeout <timeout_int>
```

Enter the number of minutes to keep the SSL session states for the segment of the SSL connection between the server and the FortiGate unit. This option appears only if `ssl-mode` is `full`.

```
ssl-server-session-state-type {both | count | disable | time}
```

Select which method the FortiGate unit should use when deciding to expire SSL sessions for the segment of the SSL connection between the server and the FortiGate unit. This option appears only if `ssl-mode` is `full`.

- **both:** Select to expire SSL session states when either `ssl-server-session-state-max` or `ssl-server-session-state-timeout` is exceeded, regardless of which occurs first.
- **count:** Select to expire SSL session states when `ssl-server-session-state-max` is exceeded.
- **disable:** Select to keep no SSL session states.
- **time:** Select to expire SSL session states when `ssl-server-session-state-timeout` is exceeded.

SSL offloading support for Internet Explorer 6

In some cases the Internet Explorer 6 web browser may be able to access real servers. To resolve this issue, disable the `ssl-send-empty-frags` option:

```
config firewall vip
  edit vip_name
    set type server-load-balance
    set server-type https
    set ssl-send-empty-frags disable
  end
```

You can disable this option if SSL acceleration will be used with an old or buggy SSL implementation that cannot properly handle empty fragments.

Selecting the cipher suites available for SSL load balancing

You can use the following command to view the complete list of cipher suites available for SSL offloading:

```
config firewall vip
edit <vip-name>
set type server-load-balance
set server-type https
set ssl-algorithm custom
config ssl-cipher-suites
edit 0
set cipher ?
```

In most configurations the matching cipher suite is automatically selected but you can limit the set of cipher suites that are available for a given SSL offloading configuration. For example, use the following command to limit an SSL load balancing configuration to use the three cipher suites that support ChaCha20 and Poly1305:

```
config firewall vip
edit <vip-name>
set type server-load-balance
set server-type https
set ssl-algorithm custom
config ssl-cipher-suites
edit 1
set cipher TLS-ECDHE-RSA-WITH-CHACHA20-POLY1305-SHA256
next
edit 2
set cipher TLS-ECDHE-ECDSA-WITH-CHACHA20-POLY1305-SHA256
next
edit 3
set cipher TLS-DHE-RSA-WITH-CHACHA20-POLY1305-SHA256
end
end
```

Disabling SSL/TLS re-negotiation

The vulnerability [CVE-2009-3555](#) affects all SSL/TLS servers that support re-negotiation. FortiOS when configured for SSL/TLS offloading is operating as a SSL/TLS server. The IETF is working on a TLS protocol change that will fix the problem identified by CVE-2009-3555 while still supporting re-negotiation. Until that protocol change is available, you can use the `ssl-client-renegotiation` option to disable support for SSL/TLS re-negotiation. The default value of this option is `allow`, which allows an SSL client to renegotiate. You can change the setting to `deny` to abort any attempts by an SSL client to renegotiate. If you select `deny` as soon as a `ClientHello` message indicating a re-negotiation is received from the client FortiOS terminates the TCP connection.

Since SSL offloading does not support requesting client certificates the only circumstance in which a re-negotiation is required is when more than 2³² bytes of data are exchanged over a single handshake. If you are sure that this volume of traffic will not occur then you can disable re-negotiation and avoid any possibility of the attack described in CVE-2009-3555.

The re-negotiation behavior can be tested using OpenSSL. The OpenSSL `s_client` application has the feature that the user can request that it do renegotiation by typing "R". For example, the following shows a successful re-negotiation against a FortiGate unit configured with a VIP for 192.168.2.100:443:

```
$ openssl s_client -connect 192.168.2.100:443
CONNECTED(00000003)
depth=1 /C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
Authority/CN=support/emailAddress=support@fortinet.com
verify error:num=19:self signed certificate in certificate chain
verify return:0
---
Certificate chain
0
s:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Fortigate/CN=FW80CM3909604325/emailAdd
ress=support@fortinet.com
i:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
Authority/CN=support/emailAddress=support@fortinet.com
1 s:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
Authority/CN=support/emailAddress=support@fortinet.com
i:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
Authority/CN=support/emailAddress=support@fortinet.com
---
Server certificate
-----BEGIN CERTIFICATE-----
---certificate not shown---
-----END CERTIFICATE-----

subject=/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Fortigate/CN=FW80CM3909604325/em
ailAddress=support@fortinet.com
issuer=/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
Authority/CN=support/emailAddress=support@fortinet.com
---
No client certificate CA names sent
---
SSL handshake has read 2370 bytes and written 316 bytes
---
New, TLSv1/SSLv3, Cipher is DHE-RSA-AES256-SHA
Server public key is 1024 bit
Compression: NONE
Expansion: NONE
SSL-Session:
    Protocol : TLSv1
    Cipher : DHE-RSA-AES256-SHA
    Session-ID:
    02781E1E368DCCE97A95396FAA82E8F740F5BBA96CF022F6FEC3597B0CC88095
    Session-ID-ctx:
    Master-Key:
    A6BBBD8477A2422D56E57C1792A4EA9C86F37D731E67D0A66E5CDB2B5C76650780C0E7F01CFF851EC44661
    86F4C48397
    Key-Arg : None
    Start Time: 1264453027
    Timeout : 300 (sec)
    Verify return code: 19 (self signed certificate in certificate
    chain)
---
```

```

GET /main.c HTTP/1.0
R
RENEGOTIATING
depth=1 /C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
Authority/CN=support/emailAddress=support@fortinet.com
verify error:num=19:self signed certificate in certificate chain
verify return:0
HTTP/1.0 200 ok
Content-type: text/plain

/*
 * Copyright (C) 2004-2007 Fortinet
 */

#include <stdio.h>
#include "vsd_ui.h"

int main(int argc, char **argv)
{
    return vsd_ui_main(argc, argv);
}
closed
$

```

The following is the same test, but this time with the VIP configuration changed to `ssl-client-renegotiation deny`:

```

$ openssl s_client -connect 192.168.2.100:443
CONNECTED(00000003)
depth=1 /C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
Authority/CN=support/emailAddress=support@fortinet.com
verify error:num=19:self signed certificate in certificate chain
verify return:0
---
Certificate chain
0
  s:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Fortigate/CN=FW80CM3909604325/emailAddress=support@fortinet.com
  i:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
  Authority/CN=support/emailAddress=support@fortinet.com
1 s:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
  Authority/CN=support/emailAddress=support@fortinet.com
  i:/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
  Authority/CN=support/emailAddress=support@fortinet.com
---
Server certificate
-----BEGIN CERTIFICATE-----
---certificate not shown---
-----END CERTIFICATE-----

  subject=/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Fortigate/CN=FW80CM3909604325/emailAddress=support@fortinet.com
  issuer=/C=US/ST=California/L=Sunnyvale/O=Fortinet/OU=Certificate
  Authority/CN=support/emailAddress=support@fortinet.com
---
No client certificate CA names sent
---

```

```

SSL handshake has read 2370 bytes and written 316 bytes
---
New, TLSv1/SSLv3, Cipher is DHE-RSA-AES256-SHA
Server public key is 1024 bit
Compression: NONE
Expansion: NONE
SSL-Session:
    Protocol : TLSv1
    Cipher : DHE-RSA-AES256-SHA
    Session-ID:
    8253331D266DDE38E4D8A04AFCA9CBDED5B1134932CE1718EED6469C1FBC7474
    Session-ID-ctx:
    Master-Key:

ED05A3EF168AF2D06A486362FE91F1D6CAA55CEFC38A3C36FB8BD74236BF2657D4701B6C1456CEB5BB5EFA
A7619EF12D
    Key-Arg : None
    Start Time: 1264452957
    Timeout : 300 (sec)
    Verify return code: 19 (self signed certificate in certificate
    chain)
---
GET /main.c HTTP/1.0
R
RENEGOTIATING

```

19916:error:1409E0E5:SSL routines:SSL3_WRITE_BYTES:ssl handshake failure:s3_pkt.c:530:

Use the following command to check the SSL stats to see that the renegotiations blocked counter is now 1:

```

diagnose firewall vip virtual-server stats ssl
ssl
  client
    connections total 0 active 0 max 0
    handshakes total 4 active 0 max 0 completed 4 abbreviated 0
    session states total 4 active 4 max 4
    cipher-suite failures 0
    embryonics total 0 active 0 max 0 terminated 0
    renegotiations blocked 1
  server
    connections total 0 active 0 max 0
    handshakes total 3 active 0 max 0 completed 2 abbreviated 1
    session states total 1 active 1 max 1
    cipher-suite failures 0
  internal error 0
  bad handshake length 0
  bad change cipher spec length 0
  pubkey too big 0
  persistence
    find 0 found 0 clash 0 addr 0 error 0

```

If the virtual server debug log is examined (diagnose debug appl vs -1) then at the point the re-negotiation is blocked there is a log:

```

vs ssl 12 handshake recv ClientHello
vs ssl 12 handshake recv 1
(0100005403014b5e056c7f573a563bebe0258c3254bbaff7046a461164f34f94f4f3d019c418000026
00390038003500160013000a00330032002f00050004001500120009001400110008000600030201000
00400230000)
vs ssl 12 client renegotiation attempted rejected, abort

```

```
vs ssl 12 closing 0 up
vs src 12 close 0 in
vs src 12 error closing
vs dst 14 error closing
vs dst 14 closed
vs ssl 14 close
vs sock 14 free
vs src 12 closed
vs ssl 12 close
vs sock 12 free
```

IP, TCP, and UDP load balancing


You can load balance all IP, TCP or UDP sessions accepted by the security policy that includes a load balancing virtual server with the type set to IP, TCP, or UDP. Traffic with destination IP and port that matches the virtual server IP and port is load balanced. For these protocol-level load balancing virtual servers you can select a load balance method and add real servers and health checking. However, you can't configure persistence, HTTP multiplexing and SSL offloading.

Example HTTP load balancing to three real web servers

In this example, a virtual web server with IP address 192.168.37.4 on the Internet, is mapped to three real web servers connected to the FortiGate unit dmz1 interface. The real servers have IP addresses 10.10.123.42, 10.10.123.43, and 10.10.123.44. The virtual server uses the **First Alive** load balancing method. The configuration also includes an HTTP health check monitor that includes a URL used by the FortiGate unit for get requests to monitor the health of the real servers.

Connections to the virtual web server at IP address 192.168.37.4 from the Internet are translated and load balanced to the real servers by the FortiGate unit. First alive load balancing directs all sessions to the first real server. The computers on the Internet are unaware of this translation and load balancing and see a single virtual server at IP address 192.168.37.4 rather than the three real servers behind the FortiGate unit.

Virtual server configuration example



GUI configuration

Use the following procedures to configure this load balancing setup from the GUI.

To add an HTTP health check monitor

In this example, the HTTP health check monitor includes the **URL** “/index.html” and the **Matched Phrase** “Fortinet products”.

1. Go to **Policy & Objects > Health Check**.
2. Select **Create New**.
3. Add an HTTP health check monitor that sends get requests to `http://<real_server_IP_address>/index.html` and searches the returned web page for the phrase “Fortinet products”.

Name	HTTP_health_chk_1
Type	HTTP
Port	80
URL	/index.html
Matched Content	Fortinet products

Interval	10 seconds
Timeout	2 seconds
Retry	3

4. Select **OK**.

To add the HTTP virtual server and the real servers

1. Go to **Policy & Objects > Virtual Servers**.
2. Select **Create New**.
3. Add an HTTP virtual server that allows users on the Internet to connect to the real servers on the internal network. In this example, the FortiGate wan1 interface is connected to the Internet.

Name	Load_Bal_VS1
Type	HTTP
Interface	wan1
Virtual Server IP	192.168.37.4 The public IP address of the web server. The virtual server IP address is usually a static IP address obtained from your ISP for your web server. This address must be a unique IP address that is not used by another host and cannot be the same as the IP address of the external interface the virtual IP will be using. However, the external IP address must be routed to the selected interface. The virtual IP address and the external IP address can be on different subnets. When you add the virtual IP, the external interface responds to ARP requests for the external IP address.
Virtual Server Port	80
Load Balance Method	First Alive
Persistence	HTTP cookie
Health Check	HTTP_health_chk_1
HTTP Multiplexing	Turn on. The FortiGate unit multiplexes multiple client into a few connections between the FortiGate unit and each real HTTP server. This can improve performance by reducing server overhead associated with establishing multiple connections.
Preserve Client IP	Turn on. The FortiGate unit preserves the IP address of the client in the <code>X-Forwarded-For</code> HTTP header.

4. Add three real servers to the virtual server. Each real server must include the IP address of a real server on the internal network.

Configuration for the first real server.

IP Address	10.10.10.42
Port	80
Max Connections	0
	Setting Max Connections to 0 means the FortiGate unit does not limit the number of connections to the real server. Since the virtual server uses First Alive load balancing you may want to limit the number of connections to each real server to limit the traffic received by each server. In this example, the Max Connections is initially set to 0 but can be adjusted later if the real servers are getting too much traffic.
Mode	Active

Configuration for the second real server.

IP Address	10.10.10.43
Port	80
Max Connections	0
Mode	Active

Configuration for the third real server.

IP Address	10.10.10.44
Port	80
Max Connections	0
Mode	Active

To add the virtual server to a security policy

Add a wan1 to dmz1 security policy that uses the virtual server so that when users on the Internet attempt to connect to the web server's IP address, packets pass through the FortiGate unit from the wan1 interface to the dmz1 interface. The virtual IP translates the destination address of these packets from the virtual server IP address to the real server IP addresses.

1. Go to **Policy & Objects > IPv4 Policy**.
2. Select **Create New**.
3. Configure the security policy:

Name	Add a name for the policy.
Incoming Interface	wan1

Outgoing Interface	dmz1
Source	all (or a more specific address)
Destination	Load_Bal_VS1
Schedule	always
Service	HTTP
Action	ACCEPT
NAT	Select this option and select Use Destination Interface Address .
Log Allowed Traffic	Select to log virtual server traffic

4. Select other security policy options as required.
5. Select **OK**.

CLI configuration

Use the following procedure to configure this load balancing setup from the CLI.

To configure HTTP load balancing

1. Use the following command to add an HTTP health check monitor that sends get requests to `http://<real_server_IP_address>/index.html` and searches the returned web page for the phrase "Fortinet products".

```
config firewall ldb-monitor
  edit HTTP_health_chk_1
    set type http
    set port 80
    set http-get /index.html
    set http-match "Fortinet products"
    set interval 10
    set timeout 2
    set retry 3
  end
```

2. Use the following command to add an HTTP virtual server that allows users on the Internet to connect to the real servers on the internal network. In this example, the FortiGate wan1 interface is connected to the Internet.

```
config firewall vip
  edit Load-Bal_VS1
    set type server-load-balance
    set server-type http
    set ldb-method first-alive
    set http-multiplex enable
    set http-ip-header enable
    set extip 192.168.37.4
    set extintf wan1
    set extport 80
    set persistence http-cookie
    set monitor HTTP_health_chk_1
    config realservers
      edit 1
        set ip 10.10.10.42
        set port 80
      end
    end
```



```

        next
        edit 2
            set ip 10.10.10.43
            set port 80
        next
        edit 3
            set ip 10.10.10.44
            set port 80
        end
    end
end

```

3. Use the following command to add a security policy that includes the load balance virtual server as the destination address.

```

config firewall policy
    edit 0
        set name <policy-name>
        set srcintf wan1
        set srcaddr all
        set dstintf dmz1
        set dstaddr Load-Bal_VS1
        set action accept
        set schedule always
        set service ALL
        set nat enable
    end

```

Configure other security policy settings as required.

Example Basic IP load balancing configuration

This example shows how to add a server load balancing virtual IP that load balances all traffic among 3 real servers. In the example the Internet is connected to `port2` and the virtual IP address of the virtual server is 192.168.20.20. The load balancing method is `weighted`. The IP addresses of the real servers are 10.10.10.1, 10.10.10.2, and 10.10.10.3. The weights for the real servers are 1, 2, and 3. The default weight is 1 and does not have to be changed for the first real server.

```

config firewall vip
    edit All_Load_Balance
        set type server-load-balance
        set server-type ip
        set extintf port2
        set extip 192.168.20.20
        set ldb-method weighted
        config realservers
            edit 1
                set ip 10.10.10.1
            next
            edit 2
                set ip 10.10.10.2
                set weight 2
            next
            edit 3
                set ip 10.10.10.3
                set weight 3
            end
        end
    end

```

end

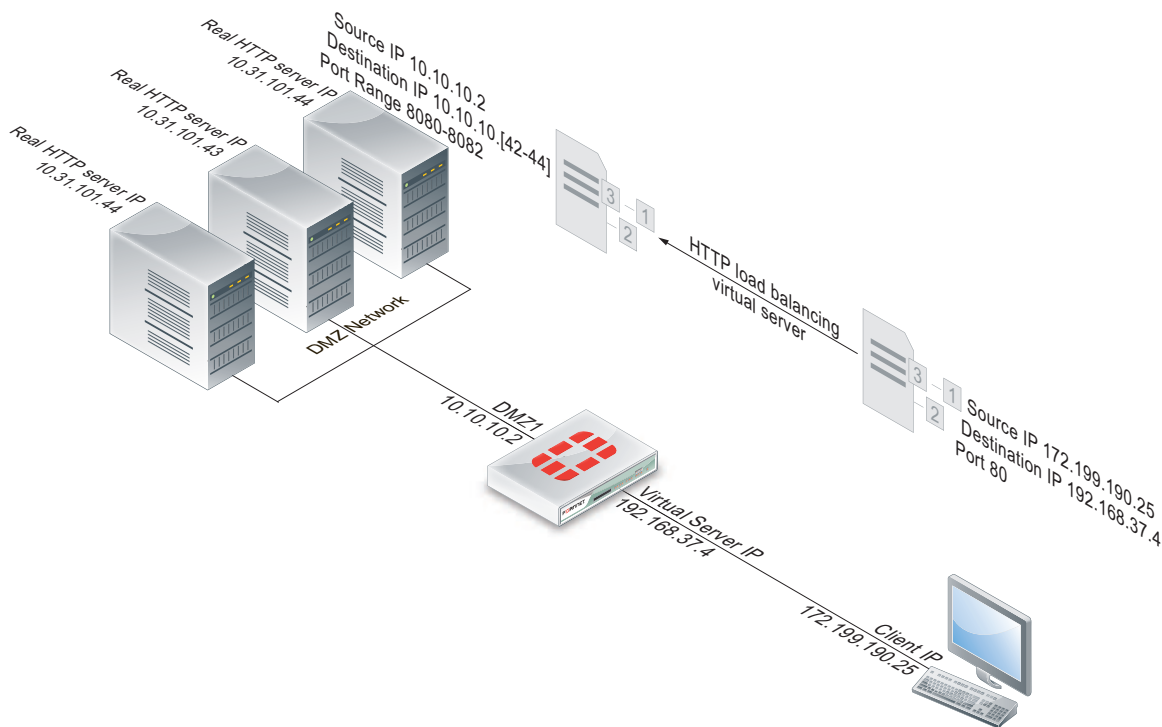
Example Adding a server load balance port forwarding virtual IP

In this example, a virtual web server with IP address 192.168.37.4 on the Internet, is mapped to three real web servers connected to the FortiGate unit dmz1 interface. The real servers have IP addresses 10.10.123.42, 10.10.123.43, and 10.10.123.44. The virtual server uses the **First Alive** load balancing method.

Each real server accepts HTTP connections on a different port number. The first real server accepts connections on port 8080, the second on port 8081, and the third on 8082. The configuration also includes an HTTP health check monitor that includes a URL used by the FortiGate unit for get requests to monitor the health of the real servers.

Connections to the virtual web server at IP address 192.168.37.4 from the Internet are translated and load balanced to the real servers by the FortiGate unit. First alive load balancing directs all sessions to the first real server. The computers on the Internet are unaware of this translation and load balancing and see a single virtual server at IP address 192.168.37.4 rather than the three real servers behind the FortiGate unit.

Server load balance virtual IP port forwarding



To complete this configuration, all of the steps would be the same as in [Example Adding a server load balance port forwarding virtual IP on page 2635](#) except for configuring the real servers.

To add the real servers to the virtual server

Use the following steps to add three real servers to the virtual server Load_Bal_VS1. These real servers cause the FortiGate unit to forward HTTP packets to the three real servers on ports 8080, 8081, and 8082.

1. Go to **Policy & Objects > Virtual Servers** and edit the **Load_Bal_VS1** virtual server.
2. Select **Create New**.
3. Add the following three real servers. Each real server must include the IP address of a real server on the internal network and have a different port number.

Configuration for the first real server.

IP Address	10.10.10.42
Port	8080
Max Connections	0
Mode	Active

Configuration for the second real server.

IP	10.10.10.43
Port	8081
Max Connections	0
Mode	Active

Configuration for the third real server.

IP	10.10.10.44
Port	8082
Max Connections	0
Mode	Active

Example Weighted load balancing configuration

This example shows how to using firewall load balancing to load balances all traffic among 3 real servers. In the example the Internet is connected to `port2` and the virtual IP address of the virtual server is 192.168.20.20. The load balancing method is `weighted`. The IP addresses of the real servers are 10.10.10.1, 10.10.10.2, and 10.10.10.3. The weights for the real servers are 1, 2, and 3.

This configuration does not include a health check monitor.

GUI configuration

Use the following procedures to configure this load balancing setup from the FortiGate GUI.

To add the HTTP virtual server

1. Go to **Policy & Objects > Virtual Servers**.
2. Select **Create New**.
3. Add an IP virtual server that allows users on the Internet to connect to the real servers on the internal network. In this example, the FortiGate port2 interface is connected to the Internet.

Name	HTTP_weghted_LB
Type	IP
Interface	port2
Virtual Server IP	192.168.20.20
Load Balance Method	Weighted

All other virtual server settings are not required or cannot be changed.

4. Under **Real Servers** select **Create New**.
5. Add three real servers. Because the **Load Balancing Method** is **Weighted**, each real server includes a weight. Servers with a higher weight receive a more sessions.

Configuration for the first real server.

IP Address	10.10.10.1
Weight	1
Max Connections	0
	Setting Max Connections to 0 means the FortiGate unit does not limit the number of connections to the real server.
Mode	Active

Configuration for the second real server.

IP Address	10.10.10.2
Weight	2
Max Connections	0
Mode	Active

Configuration for the third real server.

IP Address	10.10.10.3
Weight	3
Max Connections	0
Mode	Active

To add the virtual server to a security policy

Add a port2 to port1 security policy that uses the virtual server so that when users on the Internet attempt to connect to the web server's IP address, packets pass through the FortiGate unit from the wan1 interface to the dmz1 interface. The virtual IP translates the destination address of these packets from the virtual server IP address to the real server IP addresses.

1. Go to **Policy & Objects > IPv4 Policy**.
2. Select **Create New**.
3. Configure the security policy:

Name	Policy name
Incoming Interface	port2
Outgoing Interface	port1
Source	all (or a more specific address)
Destination	HTTP_weghted_LB
Schedule	always
Service	ALL
Action	ACCEPT
NAT	Select this option and select Use Destination Interface Address .

4. Select other security policy options as required.
5. Select **OK**.

CLI configuration

Load balancing is configured from the CLI using the `config firewall vip` command and by setting `type` to `server-load-balance`. The default weight is 1 and does not have to be changed for the first real server.

Use the following command to add the virtual server and the three weighted real servers.

```
config firewall vip
  edit HTTP_weghted_LB
    set type server-load-balance
    set server-type ip
    set extintf port2
    set extip 192.168.20.20
    set ldb-method weighted
    config realservers
      edit 1
        set ip 10.10.10.1
      next
      edit 2
        set ip 10.10.10.2
        set weight 2
      next
      edit 3
        set ip 10.10.10.3
```

```

        set weight 3
    end
end

```

Example HTTP and HTTPS persistence configuration

This example shows how to add a virtual server named **HTTP_Load_Balance** that load balances HTTP traffic using port 80 and a second virtual server named **HTTPS_Load_Balance** that load balances HTTPS traffic using port 443. The Internet is connected to port2 and the virtual IP address of the virtual server is 192.168.20.20. Both server load balancing virtual IPs load balance sessions to the same three real servers with IP addresses 10.10.10.2, 10.10.10.2, and 10.10.10.3. The real servers provide HTTP and HTTPS services.

For both virtual servers, persistence is set to **HTTP Cookie** to enable HTTP cookie persistence.

To add the HTTP and HTTPS virtual servers

1. Go to **Policy & Objects > Virtual Servers**.
2. Add the HTTP virtual server that includes HTTP Cookie persistence.

Name	HTTP_Load_Balance
Type	HTTP
Interface	port2
Virtual Server IP	192.168.20.20
Virtual Server Port	80 In this example the virtual server uses port 8080 for HTTP sessions instead of port 80.
Load Balance Method	Static
Persistence	HTTP cookie

3. Under **Real Servers** select **Create New**.
4. Add three real servers.

Configuration for the first real server.

IP Address	10.10.10.1
Port	80
Max Connections	0
Mode	Active

Configuration for the second real server.

IP Address	10.10.10.2
Port	80
Maximum Connections	0
Mode	Active

Configuration for the third real server.

IP Address	10.10.10.3
Port	80
Max Connections	0
Mode	Active

5. Select **OK**.
6. Select **Create New** to add the HTTPS virtual server that also includes HTTP Cookie persistence.

Name	HTTPS_Load_Balance
Type	HTTPS
Interface	port2
Virtual Server IP	192.168.20.20
Virtual Server Port	443
Load Balance Method	Static
Persistence	HTTP cookie

7. Under **Real Servers** select **Create New**
8. Add three real servers.

Configuration for the first real server.

IP Address	10.10.10.1
Port	443
Max Connections	0
Mode	Active

Configuration for the second real server.

IP Address	10.10.10.2
Port	443

Max Connections	0
Mode	Active

Configuration for the third real server.

IP Address	10.10.10.3
Port	443
Max Connections	0
Mode	Active

To add the virtual servers to security policies

Add a port2 to port1 security policy that uses the virtual server so that when users on the Internet attempt to connect to the web server's IP address, packets pass through the FortiGate unit from the wan1 interface to the dmz1 interface. The virtual IP translates the destination address of these packets from the virtual server IP address to the real server IP addresses.

1. Go to **Policy & Objects > IPv4 Policy**.
2. Select **Create New**.
3. Configure the HTTP security policy:

Name	Policy name.
Incoming Interface	port2
Outgoing Interface	port1
Source	all
Destination	HTTP_Load_Balance
Schedule	always
Service	HTTP
Action	ACCEPT
NAT	Select this option and select Use Destination Interface Address .

4. Select other security policy options as required.
5. Select **OK**.
6. Select **Create New**.
7. Configure the HTTP security policy:

Name	Policy name.
Incoming Interface	port2
Outgoing Interface	port1

Source	all
Destination	HTTPS_Load_Balance
Schedule	always
Service	HTTPS
Action	ACCEPT
NAT	Select this option and select Use Destination Interface Address .

8. Select other security policy options as required.

9. Select **OK**.

CLI configuration: adding persistence for a specific domain

Load balancing is configured from the CLI using the `config firewall vip` command and by setting `type` to `server-load-balance`.

For the CLI configuration, both virtual servers include setting `http-cookie-domain` to `.example.org` because HTTP cookie persistence is just required for the `example.org` domain.

First, the configuration for the HTTP virtual IP:

```
config firewall vip
  edit HTTP_Load_Balance
    set type server-load-balance
    set server-type http
    set extport 8080
    set extintf port2
    set extip 192.168.20.20
    set persistence http-cookie
    set http-cookie-domain .example.org
    config realservers
      edit 1
        set ip 10.10.10.1
      next
      edit 2
        set ip 10.10.10.2
      next
      edit 3
        set ip 10.10.10.3
    end
  end
```

Second, the configuration for the HTTPS virtual IP. In this configuration you don't have to set `extport` to 443 because `extport` is automatically set to 443 when `server-type` is set to `https`.

```
config firewall vip
  edit HTTPS_Load_Balance
    set type server-load-balance
    set server-type https
    set extport 443
    set extintf port2
    set extip 192.168.20.20
    set persistence http-cookie
    set http-cookie-domain .example.org
```

```
config realservers
  edit 1
    set ip 10.10.10.1
  next
  edit 2
    set ip 10.10.10.2
  next
  edit 3
    set ip 10.10.10.3
  end
end
```

Chapter 23 - SSL VPN

The following chapters are included in this document:

[Overview](#) provides useful general information about VPN and SSL, how the FortiGate unit implements them, and gives guidance on how to choose between SSL and IPsec.

[Basic configuration](#) explains how to configure the FortiGate unit and the web portal. Along with these configuration details, this chapter also explains how to grant unique access permissions, how to configure the SSL encryption key algorithm, and describes the SSL VPN OS Patch Check feature that allows a client with a specific OS patch to access SSL VPN services.

[The SSL VPN client](#) provides an overview of the FortiClient software required for tunnel mode, where to obtain the software, how to install it, and the configuration information required for remote users to connect to the internal network.

[The SSL VPN web portal](#) provides an overview of the SSL VPN web portal, with explanations of how to use and configure the web portal features.

[Setup examples](#) explores several configuration scenarios with step-by-step instructions. While the information provided is enough to set up the described SSL VPN configurations, these scenarios are not the only possible SSL VPN setups.

[Troubleshooting](#) provides some general maintenance and troubleshooting procedures for SSL VPNs.

What's new in FortiOS 6.0.1

The following list contains new SSL VPN features added in FortiOS 6.0.1. Click on a link to navigate to that section for further information.

- ["Visibility of SSL VPN portal SSO credentials" on page 2712](#)

What's new in FortiOS 6.0

The following list contains new SSL VPN features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- ["Tunnel Mode Client Options logic" on page 2660](#)
- ["Mac OS host check" on page 2676](#)
- ["The Launch FortiClient button" on page 2686](#)
- ["Downloading files from an SMB server in Web Mode" on page 2691](#)
- ["Split DNS support for SSL VPN portals" on page 2691](#)
- ["Automatic bookmarks for SSO credentials" on page 2694](#)
- [" HTTP header information" on page 2712](#)

Overview

As organizations have grown and become more complex, secure remote access to network resources has become critical for day-to-day operations. In addition, businesses are expected to provide clients with efficient, convenient services including knowledge bases and customer portals. Employees traveling across the country or around the world require timely and comprehensive access to network resources. As a result of the growing need for providing remote/mobile clients with easy, cost-effective and secure access to a multitude of resources, the concept of a Virtual Private Network (VPN) was developed.

SSL VPNs establish connectivity using SSL, which functions at Levels 4 - 5 (Transport and Session layers). Information is encapsulated at Levels 6 - 7 (Presentation and Application layers), and SSL VPNs communicate at the highest levels in the OSI model. SSL is not strictly a Virtual Private Network (VPN) technology that allows clients to connect to remote networks in a secure way. A VPN is a secure logical network created from physically separate networks. VPNs use encryption and other security methods to ensure that only authorized users can access the network. VPNs also ensure that the data transmitted between computers cannot be intercepted by unauthorized users. When data is encoded and transmitted over the Internet, the data is said to be sent through a "VPN tunnel". A VPN tunnel is a non-application oriented tunnel that allows the users and networks to exchange a wide range of traffic regardless of application or protocol.

The advantages of a VPN over an actual physical private network are two-fold. Rather than utilizing expensive leased lines or other infrastructure, you use the relatively inexpensive, high-bandwidth Internet. Perhaps more important though is the universal availability of the Internet. In most areas, access to the Internet is readily obtainable without any special arrangements or long wait times.

SSL (Secure Sockets Layer) as HTTPS is supported by most web browsers for exchanging sensitive information securely between a web server and a client. SSL establishes an encrypted link, ensuring that all data passed between the web server and the browser remains private and secure. SSL protection is initiated automatically when a user (client) connects to a web server that is SSL-enabled. Once the successful connection is established, the browser encrypts all the information before it leaves the computer. When the information reaches its destination, it is decrypted using a secret (private) key. Any data sent back is first encrypted, and is decrypted when it reaches the client.

FortiOS supports the SSL and TLS versions defined below:

SSL and TLS version support table

Version	RFC
SSL 2.0	RFC 6176
SSL 3.0	RFC 6101
TLS 1.0	RFC 2246
TLS 1.1	RFC 4346
TLS 1.2	RFC 5246

SSL VPN modes of operation

When a remote client connects to the FortiGate unit, the FortiGate unit authenticates the user based on username, password, and authentication domain. A successful login determines the access rights of remote users according to user group. The user group settings specify whether the connection will operate in web-only mode or tunnel mode.

Web-only mode

Web-only mode provides remote users with a fast and efficient way to access server applications from any thin client computer equipped with a web browser. Web-only mode offers true clientless network access using any web browser that has built-in SSL encryption and the Sun Java Runtime Environment (note that there is no minimum Java/JRE version requirement—any version of Java/JRE currently supported by the supplier of the Java/JRE for the operating system should work).

Support for SSL VPN web-only mode is built into FortiOS. The feature comprises of an SSL daemon running on the FortiGate unit, and a web portal, which provides users with access to network services and resources including HTTP/HTTPS, Telnet, FTP, SMB/CIFS, VNC, RDP, and SSH.

In web-only mode, the FortiGate unit acts as a secure HTTP/HTTPS gateway and authenticates remote users as members of a user group. After successful authentication, the FortiGate unit redirects the web browser to the web portal home page and the user can access the server applications behind the FortiGate unit.

When the FortiGate unit provides services in web-only mode, a secure connection between the remote client and the FortiGate unit is established through the SSL VPN security in the FortiGate unit and the SSL security in the web browser. After the connection has been established, the FortiGate unit provides access to selected services and network resources through a web portal.

FortiGate SSL VPN web portals have a 1- or 2-column page layout and portal functionality is provided through small applets called widgets. Widget windows can be moved or minimized. The controls within each widget depend on its function. There are predefined web portals and the administrator can create additional portals.

Configuring the FortiGate unit involves selecting the appropriate web portal configuration in the user group settings. These configuration settings determine which server applications can be accessed. SSL encryption is used to ensure traffic confidentiality.

The following table lists the operating systems and web browsers supported by SSL VPN web-only mode.

VPN Web-only Mode, supported operating systems and web browsers

Operating System	Web Browser
Microsoft Windows 7 SP1 (32-bit/64-bit)	<ul style="list-style-type: none">• Microsoft Internet Explorer version 11• Mozilla Firefox version 46
Microsoft Windows 8/8.1 (32-bit/64-bit)	
Microsoft Windows 10 (32-bit/64-bit)	

Operating System	Web Browser
Mac OS 10.11	<ul style="list-style-type: none"> • Safari version 9 • Chrome version 56
Linux CentOS version 6.5	<ul style="list-style-type: none"> • Mozilla Firefox version 46

Other operating systems and web browsers may function correctly, but are not supported by Fortinet.

Tunnel mode

In Tunnel mode, remote clients connect to a FortiGate unit that acts as a secure HTTP/HTTPS gateway and authenticates remote users as members of a user group.

The SSL VPN client encrypts all traffic from the remote client computer and sends it to the FortiGate unit through an SSL VPN tunnel over the HTTPS link between the user and the FortiGate unit. Another option is split tunneling, which ensures that only the traffic for the private network is sent to the SSL VPN gateway. Internet traffic is sent through the usual unencrypted route. This conserves bandwidth and alleviates bottlenecks.

When the user initiates a VPN connection with the FortiGate unit through the SSL VPN client, the FortiGate unit establishes a tunnel with the client and assigns the client a virtual IP address from a range of reserved addresses. The client uses the assigned IP address as its source address for the duration of the connection. After the tunnel has been established, the user can access the network behind the FortiGate unit.

Port forwarding mode

While tunnel mode provides a Layer 3 tunnel that users can run any application over, the user needs to install the tunnel client, and have the required administrative rights to do so. In some situations, this may not be desirable, yet the simple web mode does not provide enough flexibility for application support (for example, if you wish to use an email client that communicates with a POP3 server). The port forward mode, or proxy mode, provides this middle ground between web mode and tunnel mode.

SSL VPN port forwarding listens on local ports on the user's computer. When it receives data from a client application, the port forward module encrypts and sends the data to the FortiGate unit, which then forwards the traffic to the application server.

The port forward module is implemented with a Java applet, which is downloaded and runs on the user's computer. The applet provides the up-to-date status information such as addressing and bytes sent and received.

On the user end, the user logs into the FortiGate SSL VPN portal, and selects a port forward bookmark configured for a specific application. The bookmark defines the server address and port as well as which port to listen to on the user's computer.



The user must configure the application on the PC to point to the local proxy instead of the application server. For information on this configuration change, see the application documentation.

This mode only supports client/server applications that are using a static TCP port. It will not support client/server applications using dynamic ports or traffic over UDP.

Application support

With Citrix application servers, the server downloads an ICA configuration file to the user's PC. The client application uses this information to connect to the Citrix server. The FortiGate unit will read this file and append a SOCKS entry to set the SOCKS proxy to 'localhost'. The Citrix client will then be able to connect to the SSL VPN port forward module to provide the connection. When configuring the port forwarding module, a selection is available for Citrix servers.

For Windows Remote Desktop Connections, when selecting the RDP option, the tunnel will launch the RDP client and connect to the local loopback address after the port forward module has been initiated.

Note that the RDP/VNC web portals are **not** supported for the following platforms:

Platform	Model
FortiGate	80D, 92D, 200D, 200D-POE, 240D, 240D-POE, 600C, 800C, 1000C, 3240C, 3600C, and 5001C
FortiGate-Rugged	90D
FortiWiFi	92D

Antivirus and firewall host compatibility

The following tables list the antivirus and firewall client software packages that are supported in FortiOS.

Supported Windows XP antivirus and firewall software

Product supported	Antivirus	Firewall
Symantec Endpoint Protection V11	•	•
Kaspersky Antivirus 2009	•	
McAfee Security Center v8.1	•	•
Trend Micro Internet Security Pro	•	•
F-Secure Internet Security 2009	•	•

Supported Windows 7 32-bit and 64-bit antivirus and firewall software

Product supported	Antivirus	Firewall
CA Internet Security 2011	•	•
AVG Internet Security 2011		
F-Secure Internet Security 2011	•	•

Product supported	Antivirus	Firewall
Kaspersky Internet Security 2011	•	•
McAfee Internet Security 2011	•	•
Norton 360™ Version 4.0	•	•
Norton™ Internet Security 2011	•	•
Panda Internet Security 2011	•	•
Sophos Security Suite	•	•
Trend Micro Titanium Internet Security	•	•
ZoneAlarm Security Suite	•	•
Symantec Endpoint Protection Small Business Edition 12.0	•	•

SSL VPN conserve mode

FortiGate units perform all security profile processing in physical RAM. Since each model has a limited amount of memory, Kernel conserve mode is activated when the remaining free memory is nearly exhausted or the AV proxy has reached the maximum number of sessions it can service.

SSL VPN also has its own conserve mode. The FortiGate enters the SSL VPN conserve mode before the Kernel conserve mode in an attempt to prevent the Kernel conserve mode from triggering. During the SSL VPN conserve mode, no new SSL connections are allowed. It starts when free memory is <25% of the total memory (when the memory on the FortiGate is less than 512Mb) or <10% of the total memory (when the FortiGate has more than 512Mb built in).

To determine if the FortiGate has entered SSL VPN conserve mode - CLI

Run the following command in the CLI Console:

```
diagnose vpn ssl statistics
```

Result (showing conserve mode state in red):

```
SSLVPN statistics:
-----
Memory unit:                1
System total memory:        2118737920
System free memory:         218537984
SSLVPN memory margin:       314572800
SSLVPN state:                conserve

Max number of users:        2
Max number of tunnels:      0
Max number of connections:  13
```



```
Current number of users:      1
Current number of tunnels:    0
Current number of connections: 1
```

Traveling and security

Because SSL VPN provides a means for “on-the-go” users to dial in to the network while away from the office, you need to ensure that wherever and however they choose to dial in is secure, and not potentially compromising the corporate network.

Host check

To reinforce security, you can enable a host integrity checker to scan the remote client. The integrity checker probes the remote client computer to verify that it is safe before access is granted. Security attributes recorded on the client computer (for example, in the Windows registry, in specific files, or held in memory due to running processes) are examined and uploaded to the FortiGate unit. For more information, see [Host check on page 2674](#).

Host Check is applicable for both SSL VPN Web Mode and SSL VPN Tunnel mode.

SSL VPN and IPv6

FortiOS supports SSL VPN with IPv6 addressing, and is available for all the java applets (Telnet, VNC, RDP, and so on). IPv6 configurations for security policies and addressing include:

- Policy matching for IPv6 addresses
- Support for DNS resolving in SSL VPN
- Support IPv6 for ping
- FTP applications
- SMB

In essentially any of the following instructions, replace **IPv4** with **IPv6** to achieve the same desired results, but for IPv6 addresses and configurations.

Basic configuration

Configuring SSL VPN involves a number of configurations within FortiOS that you need to complete to make it all come together. This chapter describes the components required, and how and where to configure them to set up the FortiGate unit as an SSL VPN server. The configurations and steps are high level, to show you the procedures needed, and where to locate the options in FortiOS. For real-world examples, see [Setup examples on page 2699](#).

There are three or four key steps to configuring an SSL VPN tunnel. The first three in the points below are mandatory, while the others are optional. This chapter outlines these key steps as well as additional configurations for tighter security and monitoring.

The key steps are:

- Create user accounts and user groups for the remote clients.
([User accounts and groups on page 2651](#))
- Create a web portal to define user access to network resources.
([Configuring SSL VPN web portals on page 2656](#))
- Configure the security policies.
([Configuring security policies on page 2665](#))
- For tunnel-mode operation, add routing to ensure that client tunnel-mode packets reach the SSL VPN interface.
([Routing in tunnel mode on page 2672](#))
- Setup logging of SSL VPN activities.
([SSL VPN logs on page 2680](#))

This section contains the following information:

User accounts and groups

The first step for an SSL VPN tunnel is to add the users and user groups that will access the tunnel. You may already have users defined for other authentication-based security policies.

The user group is associated with the web portal that the user sees after logging in. You can use one policy for multiple groups, or multiple policies to handle differences between the groups such as access to different services, or different schedules.

To create a user account:

- In the web-based manager, go to **User & Device > User Definition**, and select **Create New**.
- In the CLI, use the commands in `config user local`.

All users accessing the SSL tunnel must be in a firewall user group. User names can be up to 64 characters long.

To create user groups:

- In the web-based manager, go to **User & Device > User Groups** and select **Create New**.
- In the CLI, use the commands in `config user group`.



Guest group and SSO group have been removed from `config user group` and `config vpn ssl web user-group-bookmark`.

Authentication

Remote users must be authenticated before they can request services and/or access network resources through the web portal. The authentication process can use a password defined on the FortiGate unit or optionally use established external authentication mechanisms such as RADIUS or LDAP.

To authenticate users, you can use a plain text password on the local FortiGate unit, forward authentication requests to an external RADIUS, LDAP or TACACS+ server, or utilize PKI certificates.

For information about how to create RADIUS, LDAP, TACACS+ or PKI user accounts and certificates, see the [Authentication Guide](#).

FortiOS supports LDAP password renewal notification and updates through SSL VPN. Configuration is enabled using the CLI commands:



```
config user ldap
  edit <username>
    set server <domain>
    set password-expiry-warning enable
    set password-renewal enable
  next
end
```

For more information, see the [Authentication Guide](#).

MAC host check

When a remote client attempts to log in to the portal, you can have the FortiGate unit check against the client's MAC address to ensure that only a specific computer or device is connecting to the tunnel. This can ensure better security should a password be compromised.

MAC addresses can be tied to specific portals and can be either the entire MAC address or a subset of the address. MAC host checking is configured in the CLI using the following commands:

```
conf vpn ssl web portal
  edit portal
    set mac-addr-check enable
    set mac-addr-action allow
    config mac-addr-check-rule
      edit "rule1"
        set mac-addr-list 01:01:01:01:01:01 08:00:27:d4:06:5d
        set mac-addr-mask 48
      end
    end
  end
```

IP addresses for users

After the FortiGate unit authenticates a request for a tunnel-mode connection, the FortiGate unit assigns the SSL VPN client an IP address for the session. The address is assigned from an IP Pool, which is a firewall address defining an IP address range.



Take care to prevent overlapping IP addresses. Do not assign to clients any IP addresses that are already in use on the private network. As a precaution, consider assigning IP addresses from a network that is not commonly used (for example, 10.254.254.0/24).

To set tunnel-mode client IP address range - web-based manager:

1. Go to **Policy & Objects > Addresses** and select **Create New**.
2. Enter an **Name**, for example, `SSL_VPN_tunnel_range`.
3. Select a **Type** of **IP Range**.
4. In the **Subnet/IP Range** field, enter the starting and ending IP addresses that you want to assign to SSL VPN clients, for example `10.254.254.[80-100]`.
5. In **Interface**, select **Any**.
6. Select **OK**.

To set tunnel-mode client IP address range - CLI:

If your SSL VPN tunnel range is for example 10.254.254.80 - 10.254.254.100, you could enter

```
config firewall address
  edit SSL_tunnel_users
    set type iprange
    set end-ip 10.254.254.100
    set start-ip 10.254.254.80
  end
```

Authentication of remote users

When remote users connect to the SSL VPN tunnel, they must perform authentication before being able to use the internal network resources. This can be as simple as assigning users with their own passwords, connecting to an LDAP server or using more secure options. FortiOS provides a number of options for authentication as well as security option for those connected users.

The web portal can include bookmarks to connect to internal network resources. A web (HTTP/HTTPS) bookmark can include login credentials so that the FortiGate unit automatically logs the user into the website. This means that the user logs into the SSL VPN and then does not have to enter any more credentials to visit preconfigured web sites.

Both the administrator and the end user can configure bookmarks, including SSO bookmarks. To add bookmarks as a web portal user, see [Using the Bookmarks widget on page 2692](#).

Setting the client authentication timeout

The client authentication timeout controls how long an authenticated user will remain connected. When this time expires, the system forces the remote client to authenticate again. As with the idle timeout, a shorter period of

time is more secure. The default value is 28800 seconds (8 hours). You can only modify this timeout value in the CLI.

For example, to change the authentication timeout to 18 000 seconds, enter the following commands in the CLI:

```
config vpn ssl settings
    set auth-timeout 18000
end
```

You can also set the idle timeout for the client, to define how long the user does not access the remote resources before they are logged out.

Additional timeout settings

SSL VPN timeout settings are also available to counter 'Slowloris' and 'R-U-Dead-Yet' vulnerabilities that allow remote attackers to cause a denial of service via partial HTTP requests.

The FortiGate solution involves two attributes (`http-request-header-timeout` and `http-request-body-timeout`).

CLI syntax

```
config vpn ssl settings
    set http-request-header-timeout [1-60] (seconds)
    set http-request-body-timeout [1-60] (seconds)
end
```

Allow one-time login per user

You can set the SSL VPN tunnel such that each user can only log into the tunnel one time concurrently per user per login. That is, once logged into the portal, they cannot go to another system and log in with the same credentials again.

To allow one-time login per user - web-based manager:

Go to **VPN > SSL-VPN Portals**, select a portal, and enable **Limit Users to One SSL-VPN Connection at a Time**. It is disabled by default.

To allow one-time login per user - CLI:

```
config vpn ssl web portal
    edit <portal_name>
        set limit-user-logins enable
    end
```

Strong authentication with security certificates

The FortiGate unit supports strong (two-factor) authentication through X.509 security certificates (version 1 or 3). The FortiGate unit can require clients to authenticate using a certificate, and the client can require the FortiGate unit to authenticate using a certificate.

For information about obtaining and installing certificates, see the [Authentication Guide](#).

You can select the **Require Client Certificate** option so that clients must authenticate using certificates. The client browser must have a local certificate installed, and the FortiGate unit must have the corresponding CA certificate installed.

When the remote client initiates a connection, the FortiGate unit prompts the client browser for its client-side certificate as part of the authentication process.

To require client authentication by security certificates - web-based manager:

1. Go to **VPN > SSL-VPN Settings**.
2. Select **Require Client Certificate**.
3. Select **Apply**.

To require client authentication by security certificates - CLI:

```
config vpn ssl settings
    set reqclientcert enable
end
```

If your SSL VPN clients require strong authentication, the FortiGate unit must offer a CA certificate that the client browser has installed.

In the FortiGate unit SSL VPN settings, you can select which certificate the FortiGate offers to authenticate itself. By default, the FortiGate unit offers its factory installed (Fortinet_CA_SSLProxy) certificate from Fortinet to remote clients when they connect. If you leave the default setting, a warning appears that recommends you purchase a certificate for your domain and upload it for use.

To enable FortiGate unit authentication by certificate - web-based manager:

1. Go to **VPN > SSL-VPN Settings**.
2. From the **Server Certificate** list, select the certificate that the FortiGate unit uses to identify itself to SSL VPN clients.
3. Select **Apply**.

To enable FortiGate unit authentication by certificate - CLI:

For example, to use the `example_cert` certificate

```
config vpn ssl settings
    set servercert example_cert
end
```



FortiOS will check the server certificate to verify that the certificate is valid. Only valid server certificates should be used.

NSA Suite B cryptography support

FortiOS supports the use of ECDSA Local Certificates for SSL VPN Suite B. The National Security Agency (NSA) developed Suite B algorithms in 2005 to serve as a cryptographic base for both classified and unclassified information at an interoperable level.

FortiOS allows you to import, generate, and use ECDSA certificates defined by the Suite B cryptography set. To generate ECDSA certificates, use the following command in the CLI:

```
exec vpn certificate local generate ec <certificate-name_str> <elliptic-curve-name>  
    <subject_str> [<optional_information>]
```

Configuring SSL VPN web portals

The SSL VPN portal enables remote users to access internal network resources through a secure channel using a web browser. FortiGate administrators can configure login privileges for system users as well as the network resources that are available to the users.

FortiOS supports LDAP password renewal notification and updates through SSL VPN. Configuration is enabled using the CLI commands:



```
config user ldap  
    edit <username>  
        set server <domain>  
        set password-expiry-warning enable  
        set password-renewal enable  
    next  
end
```

For more information, see the [Authentication Guide](#).

This step in the configuration of the SSL VPN tunnel sets up the infrastructure; the addressing, encryption, and certificates needed to make the initial connection to the FortiGate unit. This step is also where you configure what the remote user sees with a successful connection. The portal view defines the resources available to the remote users and the functionality they have on the network.

SSL connection configuration

To configure the basic SSL VPN settings for encryption and login options, go to **VPN > SSL-VPN Settings**.

Listen on Interface(s)	Define the interface which the FortiGate will use to listen for SSL VPN tunnel requests. This is generally your external interface.
Listen on Port	Enter the port number for HTTPS access.

Redirect port 80 to this login port

Enable to redirect the admin HTTP port to the admin HTTPS port.

There are two likely scenarios for this:

- SSL VPN is not in use, in which case the admin GUI runs on port 443 or 10443, and port 80 is redirected.
- SSL VPN runs on port 443, in which case port 80 is redirected to 443 and the admin port runs on 10443.

If the administrator chooses to run SSL VPN on port 80, the redirect option is invalid.

This can also be configured in the CLI as shown below (note that HTTPS-redirect is disabled by default):

Syntax:

```
config vpn ssl settings
    set https-redirect [enable | disable]
end
```

Restrict Access

Restrict accessibility to either **Allow access from any host** or to **Limit access to specific hosts** as desired. If selecting the latter, you must specify the hosts.

Idle Logout

Type the period of time (in seconds) that the connection can remain inactive before the user must log in again. The range is from 10 to 28800 seconds. Setting the value to 0 will disable the idle connection timeout. This setting applies to the SSL VPN session. The interface does not time out when web application sessions or tunnels are up.

Server Certificate

Select the signed server certificate to use for authentication. If you leave the default setting (Fortinet_CA_SSLProxy), the FortiGate unit offers its built-in certificate from Fortinet to remote clients when they connect. A warning appears that recommends you purchase a certificate for your domain and upload it for use.

Require Client Certificate

Select to use group certificates for authenticating remote clients. When the remote client initiates a connection, the FortiGate unit prompts the client for its client-side certificate as part of the authentication process.

For information on using PKI to provide client certificate authentication, see the [Authentication Guide](#).

Address Range

Select **Automatically assign addresses** or **Specify custom IP ranges**. The latter will allow you to select the range or subnet firewall addresses that represent IP address ranges reserved for tunnel-mode SSL VPN clients.

DNS Server

If you select **Specify**, you may enter up to two DNS servers (IPv4 or IPv6) to be provided for the use of clients.

Note: It is possible to implement a unique DNS suffix per SSL VPN portal using the CLI. Each suffix setting for each specific portal will override the `dns-suffix` setting under `config vpn ssl settings`. This is a CLI-only option, using the following syntax:

```
config vpn ssl web portal
edit <example>
set dns-suffix <string>
end
```

Specify WINS Servers

Enable to access options for entering up to two WINS servers (IPv4 or IPv6) to be provided for the use of clients.

Allow Endpoint Registration

Select so that FortiClient registers with the FortiGate unit when connecting. If you configured a registration key by going to **System > Config > Advanced**, the remote user is prompted to enter the key. This only occurs on the first connection to the FortiGate unit.

Portal configuration

The portal configuration determines what the remote user sees when they log in to the portal. Both the system administrator and the user have the ability to customize the SSL VPN portal.

To view the portals settings page, go to **VPN > SSL-VPN Portals**.

There are three pre-defined default portal configurations available:

- *full-access*
- *tunnel-access*
- *web-access*

Each portal type includes similar configuration options. Select between the different portals by double-clicking one of the default portals in the list. You can also create a custom portal by selecting the **Create New** option at the top.

Portal Setting	Description
Name	The name for the portal.
Limit Users to One SSL-VPN Connection at a Time	You can set the SSL VPN tunnel such that each user can only log into the tunnel one time concurrently per user per login. That is, once logged into the portal, they cannot go to another system and log in with the same credentials again. This option is disabled by default.
Tunnel Mode	These settings determine how tunnel mode clients are assigned IPv4 addresses.

Portal Setting	Description
Enable Split Tunneling	<p>Select so that the VPN carries only the traffic for the networks behind the FortiGate unit. The user's other traffic follows its normal route.</p> <p>If you enable split tunneling, you are required to set the Routing Address, which is the address that your corporate network is using. Traffic intended for the Routing Address will not be split from the tunnel.</p>
Source IP Pools	Select an IP Pool for users to acquire an IP address when connecting to the portal. There is always a default pool available if you do not create your own.
Tunnel Mode Client Options	<p>These options affect how the FortiClient application behaves when connected to the FortiGate VPN tunnel. When enabled, a check box for the corresponding option appears on the VPN login screen in FortiClient, and is not enabled by default.</p> <ul style="list-style-type: none"> • Allow client to save password - When enabled, if the user selects this option, their password is stored on the user's computer and will automatically populate each time they connect to the VPN. • Allow client to connect automatically - When enabled, if the user selects this option, when the FortiClient application is launched, for example after a reboot or system startup, FortiClient will automatically attempt to connect to the VPN tunnel. • Allow client to keep connections alive - When enabled, if the user selects this option, the FortiClient should try to reconnect once it detects the VPN connection is down unexpectedly (not manually disconnected by user).
Enable Web Mode	Select to enable web mode access.
Portal Message	This is a text header that appears on the top of the web portal.
Theme	Select a color styling specifically for the web portal.
Show Session Information	The Show Session Information widget displays the login name of the user, the amount of time the user has been logged in and the inbound and outbound traffic statistics.
Show Connection Launcher	Displays the Connection Launcher widget in the web portal.
Show Login History	Select to include user login history on the web portal.

Portal Setting	Description
User Bookmarks	Enable to allow users to add their own bookmarks in the web portal.
Predefined Bookmarks	Select to include bookmarks on the web portal. Bookmarks are used as links to internal network resources. When a bookmark is selected from a bookmark list, a pop-up window appears with the web page. Telnet, VNC, and RDP require a browser plugin. FTP and Samba replace the bookmarks page with an HTML file-browser.

Tunnel Mode Client Options logic

The FortiGate will check the logic of Tunnel mode VPN client options.

If `auto-connect` or `keep-alive` is enabled, the following warning message will be shown: *'save-password should be enabled if either auto-connect or keep-alive is enabled.'*

At the end of editing the portal, if either `auto-connect` or `keep-alive` is enabled and `save-password` is not enabled, the following message will be shown, and adding or editing the portal is not permitted: *'save-password should be enabled as either auto-connect or keep-alive is enabled.'*

Options to allow firewall address to be used in routing table for SSL VPN

If destination **Named Address** is set in **Network > Static Routes** and **Address Range** is set to **Automatically assign addresses** in **VPN > SSL-VPN Settings**, SSL VPN should refresh the routing table automatically.



If your network configuration does not contain a default SSL VPN portal, you might receive the error message "Input value is invalid" when you attempt to access **VPN > SSL-VPN Portals**.

To enable a default portal - CLI:

```
config vpn ssl settings
    set default-portal <full-access | tunnel-access |
    web-access>
end
```

Adding bookmarks

A web bookmark can include login credentials to automatically log the SSL VPN user into the website. When the administrator configures bookmarks, the website credentials must be the same as the user's SSL VPN credentials. Users configuring their own bookmarks can specify alternative credentials for the website.

To add a bookmark - web-based manager:

1. On the **VPN > SSL-VPN Portals** page, ensure **Enable User Bookmarks** is enabled.
2. Select **Create New** and enter the following information:

Category	Select a category, or group, to include the bookmark. If this is the first bookmark added, you will be prompted to add a category. Otherwise, select Create from the drop-down list.
Name	Enter a name for the bookmark.
Type	Select the type of link from the drop-down list. Telnet, VNC, and RDP require a browser plugin. FTP and Samba replace the bookmarks page with an HTML file-browser.
URL	Enter the IP address source.
Description	Enter a brief description of the link.
Single Sign-On	<p>Enable if you wish to use Single Sign-On (SSO) for any links that require authentication.</p> <p>When including a link using SSO, be sure to use the entire URL. For example, <code>http://10.10.1.0/login</code>, rather than just the IP address.</p>

3. Select **OK**.

For more configuration options, see [Configuring SSL VPN web portals on page 2656](#).

Personal bookmarks

The administrator has the ability to view bookmarks the remote client has added to their SSL VPN login in the bookmarks widget. This enables the administrator to monitor and, if needed, remove unwanted bookmarks that do not meet with corporate policy.

To view and maintain remote client bookmarks, go to **VPN > SSL-VPN Personal Bookmarks**.

For more information about available bookmark applications, see [Applications available in the web portal on page 2690](#)

To enable personal bookmarks:

1. Go to **System > Feature Visibility**.
2. Enable **SSL-VPN Personal Bookmark Management**.
3. Select **Apply**.

Moving and cloning bookmarks

The administrator also has the ability to move and clone personal bookmarks in the GUI and CLI.

CLI syntax

```
config vpn ssl web user-bookmark
  edit 'name'
    config bookmarks
      move bookmark1 after/before
      clone bookmark1 to
    next
  end
```

Supporting browsers without plugins (Citrix/RDPNative/Port forward) - CLI only

CLI syntax

```
config vpn ssl web user-bookmark
  edit <name>
    config bookmarks
      edit "rdpnative"
        set apptype rdpnative
        set description "rdpnative"
        set host "172.16.68.188"
        set additional-params ''
        unset full-screen-mode
        set screen-height 768
        set screen-width 1024
      next
    end
  next
end
```

Group-based SSL VPN bookmarks

The administrator can add bookmarks for groups of users. SSL VPN will only output the matched group-name entry to the client. This can only be done via the CLI.

To add group-based SSL VPN bookmarks - CLI:

```
config vpn ssl web portal
  edit "portal-name"
    set user-group-bookmark enable*/disable
  next
end
config vpn ssl web user-group-bookmark
  edit "group-name"
    config bookmark
      edit "bookmark1"
        ....
      next
    end
  next
end
```

Remote desktop bookmark creation with no password

If NLA security is chosen when creating an RDP bookmark, a username and password must be provided. However there may be instances where the user might want to use a blank password, despite being highly unrecommended. If a username is provided but the password is empty, the CLI will display a warning. See example CLI below, where the warning appears as a caution before finishing the command:

```
config vpn ssl web user-group-bookmark
  edit <group-name>
    config bookmarks
      edit <bookmark-name>
        set apptype rdp
        set host 172.16.200.121
        set security nla
        set port 3389
```

```

        set logon-user <username>
    next
end
Warning: password is empty. It might fail user authentication and remote desktop
connection would be failed.
end

```

If no username (logon-user) is specified, the following warning message will appear:

```

Please enter user name for RDP security method NLA. object set operator error, -2010
discard the setting Command fail. Return code -2010

```

SSO support for HTML5 RDP

This feature adds support for SSO from the SSL VPN portal to an RDP bookmark. If SSO is used, then the credentials used to login to SSL VPN will be automatically used when connecting to a remote RDP server.

This option is only available in CLI.

To configure SSO support for HTML5 RDP - CLI:

```

conf vpn ssl web user-bookmark
edit <name>
    config bookmarks
    edit <name>
        set apptype rdp
        set host "x.x.x.x"
        set port <value>
        set sso [disable | auto]
    next
end
next
end

```

SSL VPN Realms

You can go to **VPN > SSL-VPN Realms** and create custom login pages for your SSL VPN users. You can use this feature to customize the SSL VPN login page for your users and also to create multiple SSL VPN logins for different user groups.

In order to create a custom login page using the web-based manager, this feature must be enabled using **Feature Select**.



Before you begin, copy the default login page text to a separate text file for safe-keeping. Afterward, if needed, you can restore the text to the original version.

To configure SSL VPN Realms - web-based manager:

1. Configure a custom SSL VPN login by going to **VPN > SSL-VPN Realms** and selecting **Create New**. Users access different portals depending on the URL they enter.
2. The first option in the custom login page is to enter the path of the custom URL.
This path is appended to the address of the FortiGate unit interface to which SSL VPN users connect. The actual path for the custom login page appears beside the URL path field.
3. You can also limit the number of users that can access the custom login at any given time.

4. You can use HTML code to customize the appearance of the login page.
5. After adding the custom login, you must associate it with the users that will access the custom login. Do this by going to **VPN > SSL-VPN Settings** and adding a rule to the **Authentication/Portal Mapping** section.
6. Under **Authentication/Portal Mapping**, click **Create New** and select the user group(s) and the associated Realm.

To configure SSL VPN Realms - CLI:

```
config vpn ssl web realm
  edit <url-path>
    set login-page <content_str>
    set max-concurrent-user <int>
    set virtual-host <hostname_str>
  end
```

Where the following variables are set:

Variable	Description	Default
edit <url-path>	Enter the URL path to access the SSL-VPN login page. Do not include "http://".	No default.
login-page <content_str>	Enter replacement HTML for SSL-VPN login page.	No default.
max-concurrent-user <int>	Enter the maximum number of concurrent users allowed. Range 0-65 535. 0 means unlimited.	0
virtual-host <hostname_str>	Enter the virtual host name for this realm. Optional. Maximum length 255 characters.	No default.

Customizable FortiClient download URL

The attribute `customize-forticlient-download-url` (disabled by default) can be enabled to allow users to customize the download URL for FortiClient. This option is only available in CLI.

If enabled, two other attributes, `windows-forticlient-download-url` and `macos-forticlient-download-url`, will appear.

To configure a customizable FortiClient download URL- CLI:

```
config vpn ssl web portal
  edit <portal>
    set customize-forticlient-download-url {enable | disable}
    set windows-forticlient-download-url <custom URL for Windows>
    set macos-forticlient-download-url <custom URL for Mac OS>
  next
end
```

Disabling FortiClient download in the web portal

Use the following syntax to disable FortiClient download in the web portal.

```

config vpn ssl web portal
  edit <portal name>
    set forticlient-download disable
  next
end

```

Split DNS support

This feature allows you to specify which domains will be resolved by the DNS server specified by the VPN while all other domains will be resolved by the locally specified DNS. This is useful in both Enterprise and MSP scenarios (when hosting multiple SSL VPN portals). This option is only available in CLI.

To configure split DNS support - CLI:

```

config vpn ssl web portal
  edit <name>
    config split-dns-domains
      edit 1
        set domains "abc.com, cde.com"
        set dns-server1 192.168.1.1
        set dns-server2 192.168.1.2
        set ipv6-dns-server1 2000:2:3:4::5
        set ipv6-dns-server2 2000:2:3:4::6
      next
    ...
  end
end

```

Configuring security policies

You will need at least one SSL VPN security policy. This is an identity-based policy that authenticates users and enables them to access the SSL VPN web portal. The SSL VPN user groups named in the policy determine who can authenticate and which web portal they will use. From the web portal, users can access protected resources or download the SSL VPN tunnel client application.

This section contains the procedures needed to configure security policies for web-only mode operation and tunnel-mode operation. These procedures assume that you have already completed the procedures outlined in [Configuring security policies on page 2665](#).

If you will provide tunnel mode access, you will need a second security policy—an ACCEPT tunnel mode policy to permit traffic to flow between the SSL VPN tunnel and the protected networks.

Firewall addresses

Before you can create security policies, you need to define the firewall addresses you will use in those policies. For both web-only and tunnel mode operation, you need to create firewall addresses for all of the destination networks and servers to which the SSL VPN client will be able to connect.

For tunnel mode, you will already have defined firewall addresses for the IP address ranges that the FortiGate unit will assign to SSL VPN clients.

The source address for your SSL VPN security policies will be the predefined “all” address. Both the address and the netmask are 0.0.0.0. The “all” address is used because VPN clients will be connecting from various addresses, not just one or two known networks. For improved security, if clients will be connecting from one or

two known locations you should configure firewall addresses for those locations, instead of using the “all” address.

To create a firewall address, in the web-based manager, go to **Policy & Objects > Objects > Addresses**, and select **Create New**.

Create an SSL VPN security policy

At minimum, you need one SSL VPN security policy to authenticate users and provide access to the protected networks. You will need additional security policies only if you have multiple web portals that provide access to different resources. You can use one policy for multiple groups, or multiple policies to handle differences between the groups such as access to different services, or different schedules.

The SSL VPN security policy specifies:

- The incoming interface that corresponds to the ssl.root interface.
- The SSL VPN user groups that can use the security policy.
- The times (schedule) and types of services that users can access.
- The UTM features and logging that are applied to the connection.



Do not use ALL as the destination address. If you do, you will see the “Destination address of Split Tunneling policy is invalid” error when you enable Split Tunneling.

To create an SSL-VPN security policy - web-based manager:

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information:

Incoming Interface	Select the virtual SSL VPN interface, such as ssl.root .
Outgoing Interface	Select the FortiGate network interface that connects to the protected network.
Source	Select to allow access only to holders of a (shared) group certificate. The holders of the group certificate must be members of an SSL VPN user group, and the name of that user group must be present in the Allowed field. See Configuring security policies on page 2665 .
Source User Group	SSL VPN
Destination Address	<p>Select the firewall address you created that represents the networks and servers to which the SSL VPN clients will connect.</p> <p>If you want to associate multiple firewall addresses or address groups with the Destination Interface/Zone, from Destination Address, select the plus symbol. In the dialog box, move the firewall addresses or address groups from the Available Addresses section to the Members section, then select OK.</p>
Schedule	Select always .

Service	Select services in the left list and use the right arrow button to move them to the right list. Select the ALL service to allow the user group access to all services.
Action	Select Accept .

Your identity-based policies are listed in the security policy table. The FortiGate unit searches the table from the top down to find a policy to match the client's user group. Using the move icon in each row, you can change the order of the policies in the table to ensure the best policy will be matched first. You can also use the icons to edit or delete policies. Furthermore, you can drag and drop policies in the policy list to rearrange their order.

To create an SSL VPN security policy - CLI:

Create the SSL VPN security policy by entering the following CLI commands.

```
config firewall policy
  edit <id>
    set srcintf ssl.root(sslvpn tunnel interface)
    set dstintf port2
    set srcaddr all
    set dstaddr OfficeLAN
    set nat enable
    set groups <name>
  end
```

Create a tunnel mode security policy

If your SSL VPN will provide tunnel mode operation, you need to create a security policy to enable traffic to pass between the SSL VPN virtual interface and the protected networks. This is in addition to the SSL VPN security policy that you created in the preceding section.

The SSL VPN virtual interface is the FortiGate unit end of the SSL tunnel that connects to the remote client. It is named `ssl.<vdom_name>`. In the root VDOM, for example, it is named `ssl.root`. If VDOMs are not enabled on your FortiGate unit, the SSL VPN virtual interface is also named `ssl.root`.

To configure the tunnel mode security policy - web-based manager:

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information and select **OK**.

Incoming Interface	Select the virtual SSL VPN interface, such as ssl.root .
Outgoing Interface	Select the FortiGate network interface that connects to the protected network.
Source Address	Select the firewall address you created that represents the IP address range assigned to SSL VPN clients, such as <code>SSL_VPN_tunnel_users</code> .
Source User(s)	Select to allow access only to holders of a (shared) group certificate. The holders of the group certificate must be members of an SSL VPN user group, and the name of that user group must be present in the Allowed field. See Configuring security policies on page 2665 .

Destination Address	Select the firewall address that represents the networks and servers to which the SSL VPN clients will connect. To select multiple firewall addresses or address groups, select the plus sign next to the drop-down list.
Schedule	Select always .
Service	Select services in the left list and use the right arrow button to move them to the right list. Select the ALL service to allow the user group access to all services.
Action	Select Accept .
Enable NAT	Select Enable NAT. (Optional)

To configure the tunnel mode security policy - CLI:

```
config firewall policy
  edit <id>
    set srcintf ssl.root(sslvpn tunnel interface)
    set dstintf <dst_interface_name>
    set srcaddr <tunnel_ip_address>
    set dstaddr <protected_network_address_name>
    set schedule always
    set service ALL
    set nat enable
    set groups <name>
  end
```

This policy enables the SSL VPN client to initiate communication with hosts on the protected network. If you want to enable hosts on the protected network to initiate communication with the SSL VPN client, you should create another Accept policy like the preceding one but with the source and destination settings reversed.

You must also add a static route for tunnel mode operation.

Routing for tunnel mode

If your SSL VPN operates in tunnel mode, you must add a static route so that replies from the protected network can reach the remote SSL VPN client.

To add the tunnel mode route - web-based manager:

1. Go to **Network > Static Routes** and select **Create New**.

For low-end FortiGate units, go to **System > Network > Routing** and select **Create New**.

2. Enter the **Destination IP/Mask** of the tunnel IP address that you assigned to the users of the web portal.
3. Select the SSL VPN virtual interface for the **Device**.
4. Select **OK**.

To add the tunnel mode route - CLI:

If you assigned 10.11.254.0/24 as the tunnel IP range, you would enter:

```
config router static
```

```

edit <id>
    set device ssl.root
    set dst 10.11.254.0/24
end

```

Split tunnel Internet browsing policy

With split tunneling disabled, all of the SSL VPN client's requests are sent through the SSL VPN tunnel. But the tunnel mode security policy provides access only to the protected networks behind the FortiGate unit. Clients will receive no response if they attempt to access Internet resources. You can enable clients to connect to the Internet through the FortiGate unit using a split tunnel Internet browsing policy.

To add an Internet browsing policy:

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information and select **OK**.

Incoming Interface	Select the virtual SSL VPN interface (ssl.root , for example).
Outgoing Interface	Select the FortiGate network interface that connects to the Internet.
Source	Select the firewall address you created that represents the IP address range assigned to SSL VPN clients.
Source User Group	SSL VPN
Destination Address	Select All .
Action	Select Accept .
Enable NAT	Select Enable .

To configure the Internet browsing security policy - CLI:

To enable browsing the Internet through port1, you would enter:

```

config firewall policy
edit 0
    set srcintf ssl.root
    set dstintf port1
    set srcaddr SSL_tunnel_users
    set dstaddr all
    set schedule always
    set service ALL
    set nat enable
    set groups <name>
end

```

Enabling a connection to an IPsec VPN

You might want to provide your SSL VPN clients access to another network, such as a branch office, that is connected by an IPsec VPN. To do this, you need only to add the appropriate security policy. For information about route-based and policy-based IPsec VPNs, see the [IPsec VPN Guide](#).

Route-based connection

To configure interconnection with a route-based IPsec VPN - web-based manager:

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information and select **OK**.

Incoming Interface	Select the virtual SSL VPN interface (ssl.root , for example).
Outgoing Interface	Select the virtual IPsec interface for your IPsec VPN.
Source	Select the firewall address that represents the IP address range assigned to SSL VPN clients.
Source User Group	SSL VPN
Destination Address	Select the address of the IPsec VPN remote protected subnet.
Action	Select ACCEPT .
Enable NAT	Enable.

To configure interconnection with a route-based IPsec VPN - CLI:

If, for example, you want to enable SSL VPN users to connect to the private network (address name OfficeAnet) through the toOfficeA IPsec VPN, you would enter:

```
config firewall policy
  edit 0
    set srcintf ssl.root
    set dstintf toOfficeA
    set srcaddr SSL_tunnel_users
    set dstaddr OfficeAnet
    set action accept
    set nat enable
    set schedule always
    set service ALL
    set groups <name>
  end
```

Policy-based connection

To configure interconnection with a policy-based IPsec VPN - web-based manager:

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information and select **OK**.

Incoming Interface	Select the virtual SSL VPN interface (ssl.root , for example).
Outgoing Interface	Select the FortiGate network interface that connects to the Internet.
Source	Select the firewall address that represents the IP address range assigned to SSL VPN clients.

Source User Group	SSL VPN
Destination Address	Select the address of the IPsec VPN remote protected subnet.

3. Configure inbound NAT from the CLI:

```
config firewall policy
  edit 0
    set natinbound enable
    set groups <name>
  end
```

To configure interconnection with a policy-based IPsec VPN - CLI:

If, for example, you want to enable SSL VPN users to connect to the private network (address name OfficeAnet) through the OfficeA IPsec VPN, you would enter:

```
config firewall policy
  edit 0
    set srcintf ssl.root
    set dstintf port1
    set srcaddr SSL_tunnel_users
    set dstaddr OfficeAnet
    set action ipsec
    set schedule always
    set service ALL
    set inbound enable
    set outbound enable
    set natinbound enable
    set vpntunnel OfficeA
    set groups <name>
  end
```

In this example, port1 is connected to the Internet.

Configuring encryption key algorithms

The FortiGate unit supports a range of cryptographic cipher suites to match the capabilities of various web browsers. The web browser and the FortiGate unit negotiate a cipher suite before any information (for example, a user name and password) is transmitted over the SSL link. You can only configure encryption key algorithms for SSL VPN in the CLI.

To configure encryption key algorithms - CLI:

Use the following CLI command,

```
config vpn ssl settings
  set algorithm <cipher_suite>
end
```

where one of the following variables replaces *<cipher_suite>*:

Variable	Description
low	Use any cipher suite; AES, 3DES, RC4, DES, or ChaCha.
medium	Use a 128-bit or greater cipher suite; AES, 3DES, RC4, or ChaCha.
high	Use a cipher suite greater than 128 bits; AES or ChaCha.

Note that the `algorithm <cipher_suite>` syntax is only available when the `sslvpn-enable` attribute is set to **enable**.

Controlling the use of specific cipher suites

Administrators can ban the use of specific cipher suites in the CLI for SSL VPN, so PCI-DSS (Payment Card Industry Data Security Standard) certification can be met.

To ban the use of specific cipher suites for SSL VPN - CLI:

```
config vpn ssl settings
  set banned-cipher [RSA | DH | DHE | ECDH | ECDHE | DSS | ECDSA | AES | AESGCM |
    CAMELLIA | 3DES | SHA1 | SHA256 | SHA384 | STATIC]
```

Additional configuration options

Beyond the basics of setting up the SSL VPN, you can configure a number of other options that can help to ensure your internal network is secure and can limit the possibility of attacks and viruses entering the network from an outside source.

Routing in tunnel mode

If you are creating a SSL VPN connection in tunnel mode, you need to add a static route so that replies from the protected network can reach the remote SSL VPN client.

To add the tunnel mode route - web-based manager:

1. Go to **Network > Static Routes** and select **Create New**.
2. Enter the **Destination IP/Mask** of the tunnel IP address that you assigned to the users of the web portal.
3. Select the SSL VPN virtual interface for the **Device**.
4. Select **OK**.

To add the tunnel mode route - CLI:

If you assigned 10.11.254.0/24 as the tunnel IP range, you would enter:

```
config router static
  edit <id>
    set device ssl.root
    set dst 10.11.254.0/24
  end
```

Changing the port number for web portal connections

You can specify a different TCP port number for users to access the web portal login page through the HTTPS link. By default, the port number is 443 and users can access the web portal login page using the following default URL:

```
https://<FortiGate_IP_address>:443/remote/login
```

where <FortiGate_IP_address> is the IP address of the FortiGate interface that accepts connections from remote users.

To change the SSL VPN port - web-based manager:

1. If **Current VDOM** appears at the bottom left of the screen, select **Global** from the list of VDOMs.
2. Go to **VPN > SSL-VPN Settings**.
3. Type an unused port number in the **Listen on Port** field and select **Apply**.

To change the SSL VPN port - CLI:

This is a global setting. For example, to set the SSL VPN port to 10443, enter the following:

```
config vpn ssl settings
    set port 10443
end
```

HTTP to HTTPS redirect support

The admin HTTP port can be redirected to the admin HTTPS port. This is enabled in **VPN > SSL-VPN Settings** using the option **Redirect port 80 to this login port**.

There are two likely scenarios for this:

- SSL VPN is not in use, in which case the admin GUI runs on port 443 or 10443, and port 80 is redirected.
- SSL VPN runs on port 443, in which case port 80 is redirected to 443 and the admin port runs on 10443.

If the administrator chooses to run SSL VPN on port 80, the redirect option is invalid.

This can also be configured in the CLI as described below:

To redirect HTTP to HTTPS port - CLI:

```
config vpn ssl settings
    set https-redirect [enable | disable] (default: disabled)
end
```

SSL offloading

To configure SSL offloading, which allows or denies client renegotiation, you must use the CLI. This helps to resolve the issues that affect all SSL and TLS servers that support renegotiation, identified by the Common Vulnerabilities and Exposures system in CVE-2009-3555. The SSL offloading renegotiation feature is considered a workaround until the IETF permanently resolves the issue.

The CLI command is `ssl-client-renegotiation` and is found under the `config firewall vip` syntax.

Host check

When you enable AV, FW, or AV-FW host checking in the web portal Security Control settings, each client is checked for security software that is recognized by the Windows Security Center. As an alternative, you can create a custom host check that looks for security software selected from the Host Check list. For more information, see [Additional configuration options on page 2672](#).

The Host Check list includes default entries for many security software products.



Host integrity checking is only possible with client computers running Microsoft Windows platforms.

To configure host checking - CLI:

To configure the full-access portal to check for AV and firewall software on client Windows computers, you would enter the following:

```
config vpn ssl web portal
  edit full-access
    set host-check av-fw
  end
```

To configure the full-access portal to perform a custom host check for FortiClient Host Security AV and firewall software, you would enter the following:

```
config vpn ssl web portal
  edit full-access
    set host-check custom
    set host-check-policy FortiClient-AV FortiClient-FW
  end
```

Replacing the host check error message

You can add your own host security check error message using either the web-based manager or the CLI. The default message reads: “Your PC does not meet the host checking requirements set by the firewall. Please check that your OS version or antivirus and firewall applications are installed and running properly or you have the right network interface.”

To replace the host check error message - web-based manager:

1. Navigate to **System > Replacement Messages** and select **Extended View** in the upper right corner.
2. Scroll down to **SSL VPN** and select **Hostcheck Error Message**.
3. Edit the text in the right-hand column below and select **Save**.
If you are unhappy with the new message, you can restore the message to its default by selecting **Restore Default** instead of **Save**.

To replace the host check error message - CLI:

Configure the host check error message using the following command.

```
config system replacemsg sslvpn hostcheck-error
```

Creating a custom host check list

You can add your own software requirements to the host check list using the CLI. Host integrity checking is only possible with client computers running Microsoft Windows platforms. Enter the following commands:

```
config vpn ssl web host-check-software
  edit <software_name>
    set guid <guid_value>
    set type <av | fw>
    set version <version_number>
  end
```

If known, enter the Globally Unique Identifier (GUID) for the host check application. Windows uses GUIDs to identify applications in the Windows Registry. The GUID can be found in the Windows registry in the HKEY_CLASSES_ROOT section.

To obtain the exact versioning, in Windows, right-click on the .EXE file of the application and select **Properties**, then select the **Version** tab.

Example Tunnel Mode Host Check - Registry Key Check

- Check to see if a required registry key is present:

```
config vpn ssl web host-check-software
  edit <computer_name>
    config check-item-list
      edit 1
        set target "HKEY_LOCAL_MACHINE\\SYSTEM\\CurrentControlSet\\Control\\ComputerName\\ActiveComputerName:ComputerName=WINXP32SP3B62"
        set type registry <<<-----
      next
    end
  next
end
```

Example Tunnel Mode Host Check - Application Running Check

- Check to see if a required application is installed and/or running:

```
config vpn ssl web host-check-software
  edit "calc"
    config check-item-list
      edit 1
        set target "calc.exe"
        set type process <<<-----
      next
    end
  next
end
```

Example Tunnel Mode Host Check - File Check

- Check to see if a specific file exists at a specific location:

```
config vpn ssl web host-check-software
  edit "putty"
    config check-item-list
      edit 1
```

```

        set target "C:\\software\\putty.txt"
        set md5s <ENC>
    next
end
next
end

```

Mac OS host check

This feature provides the host check function for Mac OS to SSL VPN. The following Mac OS hosts can be allowed, denied, or checked:

- macos-high-sierra-10.13
- macos-sierra-10.12
- os-x-el-capitan-10.11
- os-x-yosemite-10.10
- os-x-mavericks-10.9

The `os-type` option is available under `vpn ssl web host-check-software`; if `os-type` is `macos`, then `type`, `version` and `guid` are hidden. Furthermore, `type` in `check-item-list` can only be set to `file` or `process`.

SSL VPN Host check support is currently not planned for iOS or Android.

To configure Mac OS host check - CLI:

```

config vpn ssl web portal
  edit <name>
    set os-check enable
    config os-check-list macos-high-sierra-10.13
      set action {allow | deny | check-up-to-date}
      set tolerance <value>
      set latest-patch-level <value>
    next
  end
next
end

config vpn ssl web host-check-software
  edit <name>
    set os-type macos
    config check-item-list
      edit <name>
        set type process
        set target "calc.exe"
      next
    end
  [...]
next
end

```

Configuring virtual desktop

Available for 32-bit Windows XP, Windows Vista, and Windows 7 client PCs, the virtual desktop feature completely isolates the SSL VPN session from the client computer's desktop environment. All data is encrypted,

including cached user credentials, browser history, cookies, temporary files, and user files created during the session. When the SSL VPN session ends normally, the files are deleted. If the session ends due to a malfunction, files might remain, but they are encrypted so that the information is protected.

When the user starts an SSL VPN session that has virtual desktop enabled, the virtual desktop replaces the user's normal desktop. When the virtual desktop exits, the user's normal desktop is restored.

Virtual desktop requires the Fortinet cache cleaner plugin. If the plugin is not present, it automatically downloads to the client computer.

To enable virtual desktop :

To enable virtual desktop on the full-access portal and apply the application control list 'List1', for example, you would enter:

```
config vpn ssl web portal
  edit full-access
    set virtual-desktop enable
    set virtual-desktop-app-list List1
  end
```

Configuring virtual desktop application control

You can control which applications users can run on their virtual desktop. To do this, you create an Application Control List of either allowed or blocked applications. When you configure the web portal, you select the list to use.

Configure the application control list in the CLI.

To create an Application Control List - CLI:

If you want to add 'BannedApp' to 'List1', a list of blocked applications, you would enter:

```
config vpn ssl web virtual-desktop-app-list
  edit "List1"
    set action block
    config apps
      edit "BannedApp"
        set md5s "06321103A343B04DF9283B80D1E00F6B"
      end
    end
  end
```

Configuring client OS Check

The SSLVPN client OS Check feature can determine if clients are running the Windows 2000, Windows XP, Windows Vista, Windows 7, or Windows 10 operating system. You can configure the OS Check to do any of the following:

- Allow the client access.
- Allow the client access only if the operating system has been updated to a specified patch (service pack) version.
- Deny the client access.

The OS Check has no effect on clients running other operating systems.

The Windows patch check enables you to define the minimum Windows version and patch level allowed when connecting to the SSL VPN portal. When the user attempts to connect to the web portal, FortiOS performs a

query on the version of Windows the user has installed. If it does not match the minimum requirement, the connection is denied. The Windows patch check is configured in the CLI.

To specify the acceptable patch level, you set the `latest-patch-level` and the `tolerance`. The lowest acceptable patch level is `latest-patch-level` minus `tolerance`. In this case, `latest-patch-level` is 3 and `tolerance` is 1, so 2 is the lowest acceptable patch level.

To configure OS Check:

OS Check is configurable only in the CLI.

```
config vpn ssl web portal
  edit <portal_name>
    set os-check enable
    config os-check-list [windows-2000 | windows-xp | windows-vista | windows-7 |
      windows-10]
      set action [allow | check-up-to-date | deny]
      set latest-patch-level [disable | 0 - 255]
      set tolerance <tolerance_num>
    end
  end
```

Host check for Windows firewall

The Windows built-in firewall does not have a GUID in `root\securitycenter` or `root\securitycenter2`, but you can use a registry value to detect the firewall status.

If Windows firewall is on, the following registry value will be set to 1:

- **KeyName:** `HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\SharedAccess\Parameters\FirewallPolicy\StandardProfile`
- **ValueName:** `EnableFirewall`

In FortiOS, use the `registry-value-check` feature to define the Windows Firewall software by entering the following in the CLI:

```
config vpn ssl web host-check-software
  edit "Microsoft-Windows-Firewall"
    config check-item-list
      edit 1
        set target
          "HKLM\\SYSTEM\\CurrentControlSet\\Services\\SharedAccess\\Parameters\\Firew
            allPolicy\\StandardProfile:EnableFirewall==1"
        set type registry
      next
      edit 2
        set target
          "HKLM\\SYSTEM\\CurrentControlSet\\Services\\SharedAccess\\Parameters\\Firew
            allPolicy\\PublicProfile:EnableFirewall==1"
        set type registry
      next
      edit 3
        set target
          "HKLM\\SYSTEM\\CurrentControlSet\\Services\\SharedAccess\\Parameters\\Firew
            allPolicy\\DomainProfile:EnableFirewall==1"
        set type registry
      next
    end
```

```
set type fw
next
set host-check custom
set host-check-policy Microsoft-Windows-Firewall
```

Adding WINS and DNS services for clients

You can specify the WINS or DNS servers that are made available to SSL-VPN clients.

DNS servers provide the IP addresses that browsers need to access web sites. For Internet sites, you can specify the DNS server that your FortiGate unit uses. If SSL VPN users will access intranet sites using URLs, you need to provide them access to the intranet's DNS server. You specify a primary and a secondary DNS server.

A WINS server provides IP addresses for named servers in a Windows domain. If SSL VPN users will access a Windows network, you need to provide them access to the domain WINS server. You specify a primary and a secondary WINS server.

To specify WINS and DNS services for clients - web-based manager:

1. Go to **VPN > SSL-VPN Settings**.
2. Next to **DNS Server** select **Specify**.
3. Enter the IP addresses of DNS servers in the **DNS Server** fields as needed. Fields are available for both IPv4 and IPv6 addresses.
4. Select **Specify WINS Servers**, and enter the IP addresses of WINS servers in the **WINS Server** fields as needed. Fields are available for both IPv4 and IPv6 addresses.
5. Select **Apply**.

To specify WINS and DNS services for clients - CLI:

```
config vpn ssl settings
  set dns-server1 <address_ipv4>
  set dns-server2 <address_ipv4>
  set wins-server1 <address_ipv4>
  set wins-server2 <address_ipv4>
end
```

Idle timeout

The idle timeout setting controls how long the connection can remain idle before the system forces the remote user to log in again. For security, keep the default value of 5000 seconds or less. Set the timeout value to 0 to disable idle timeouts.

To set the idle timeout - web-based manager:

1. Go to **VPN > SSL-VPN Settings** and enable **Idle Logout**.
2. In the **Inactive For** field, enter the timeout value.
The valid range is from 10 to 28800 seconds.
3. Select **Apply**.

To set the idle timeout - CLI:

```
config vpn ssl settings
  set idle-timeout <seconds_int>
```

```
end
```

Login timeout

With long network latency, the FortiGate can timeout the client before it can finish negotiation processes, such as DNS lookup and time to enter a token. Two CLI commands under `config vpn ssl settings` allow the login timeout to be configured, replacing the previous hard timeout value. The second command can be used to set the SSL VPN maximum DTLS hello timeout.

CLI syntax

```
config vpn ssl settings
  edit <example>
    set login-timeout [10-180] Default is 30 seconds.
    set dtls-hello-timeout [10-60] Default is 10 seconds.
end
```

Login failure limit

The following CLI allows the administrator to configure the number of times wrong credentials are allowed before the SSL VPN server blocks an IP address, and also how long the block would last.

CLI syntax

```
config vpn ssl settings
  set login-attempt-limit [0-10] Default is 2.
  set login-block-time [0-86400] Default is 60 seconds.
end
```

SSL VPN logs

Logging is available for SSL VPN traffic so you can monitor users connected to the FortiGate unit and their activity. For more information on configuring logs on the FortiGate unit, see the [Logging and Reporting Guide](#).

To enable logging of SSL VPN events - web-based manager:

1. Go to **Log & Report > Log Settings**.
2. Enable **Event Logging**, and select **VPN activity event**.
3. Select **Apply**.

To view the SSL VPN log data, in the web-based manager, go to **Log & Report** and select either the **Event Log** or **Traffic Log**.

In event log entries, look for the sub-types “sslvpn-session” and “sslvpn-user”.

For information about how to interpret log messages, see the [FortiGate Log Message Reference](#).

Monitoring active SSL VPN sessions

You can go to **User & Device > Monitor** to view a list of active SSL VPN sessions. The list displays the user name of the remote user, the IP address of the remote client, and the time the connection was made. You can also see which services are being provided, and delete an active web session from the FortiGate unit.

To monitor SSL VPNs - web-based manager:

To view the list of active SSL VPN sessions, go to **Monitor > SSL-VPN Monitor**.

When a tunnel-mode user is connected, the **Description** field displays the IP address that the FortiGate unit assigned to the remote host.

If required, you can end a session/connection by selecting its checkbox and then clicking the **Delete** icon.

Importing and using a CA-signed SSL certificate

Use the following set of instructions to import a CA-signed SSL certificate and configure an SSL VPN using that certificate.

Import the signed certificate into your FortiGate device

1. Unzip the file downloaded from the CA.
There should be two .CRT files: a CA certificate with bundle in the file name, and a local certificate.
2. Log in to your FortiGate unit and browse to **System > Certificates**.
3. Select **Create New > Local Certificate** to import the local certificate.
The status of the certificate will change from PENDING to OK.
4. Import the CA certificate by selecting **Import > CA Certificate**.
It will be listed in the CA Certificates section of the certificates list. You can now configure SSL VPN using the signed certificate.

Configure your FortiGate device to use the signed certificate

1. Log in to your FortiGate unit and browse to **VPN > SSL-VPN Settings**.
2. In the **Connection Settings** section, locate the **Server Certificate** field.
3. Select the new certificate from the drop-down menu.
4. Select **Apply** to configure SSL VPN to use the new certificate.

Implement post-authentication CSRF protection in SSL VPN web mode

This attribute can enable/disable verification of a referrer in the HTTP request header in order to prevent a Cross-Site Request Forgery attack.

CLI Syntax

```
config vpn ssl settings
    set check-referer [enable|disable]
end
```

DTLS support

The Datagram Transport Layer Security (DTLS) protocol is supported for SSL VPN connections. DTLS allows datagram-based applications to communicate in a way that prevents eavesdropping, tampering, or message forgery. It can also be used to improve upload/download throughput. It is similar to the Transport Layer Security (TLS) protocol.

DTLS support can be enabled in the CLI as described below.

CLI Syntax

```
config vpn ssl settings
    set dtls-tunnel [enable | disable] (default: enabled)
end
```


Allow firewall address to be used in routing table for SSL VPN

If destination **Named Address** is set in **Network > Static Routes** and **Address Range is set to Automatically assign addresses** is enabled in **VPN > SSL-VPN Settings**, SSL VPN should refresh the routing table automatically.

To view the routes in the routing table, go to **Monitor > Routing Monitor**.

WAN link load balancing

You can set `virtual-wan-link` as the destination interface in a firewall policy (when SSL VPN is the source interface) for WAN link load balancing. This allows logging into a FortiGate via SSL VPN for traffic inspection and then have outbound traffic load balanced by WAN link load balancing.

CLI syntax

```
config firewall policy
  edit <example>
    set dstintf virtual-wan-link
  end
```

The SSL VPN client

The remote client connects to the SSL VPN tunnel in various ways, depending on the VPN configuration.

- Tunnel mode establishes a connection to the remote protected network that any application can use. If the client computer runs Microsoft Windows, they can download the tunnel mode client from the web portal. If the client computer runs Linux or Mac OS X, the user needs to download the tunnel mode client application from the Fortinet Support web site. See the Release Notes for your FortiOS firmware for the specific operating system versions that are supported.
- The virtual desktop application creates a virtual desktop on a user's PC and monitors the data read/write activity of the web browser running inside the virtual desktop. When the application starts, it presents a 'virtual desktop' to the user. The user starts the web browser from within the virtual desktop and connects to the SSL VPN web portal. The browser file/directory operation is redirected to a new location, and the data is encrypted before it is written to the local disk. When the virtual desktop application exits normally, all the data written to the disk is removed. If the session terminates abnormally (power loss, system failure, etc.), the data left behind is encrypted and unusable to the user. The next time you start the virtual desktop, the encrypted data is removed.



The SSL VPN standalone client installer for Windows is no longer supported in FortiOS 5.4. Users should use the FortiClient installer with the "VPN Only" option instead.

FortiClient

Remote users can use the FortiClient software to initiate an SSL VPN tunnel to connect to the internal network. FortiClient uses local port TCP 1024 to initiate an SSL encrypted connection to the FortiGate unit, on port TCP 443. When connecting using FortiClient, the FortiGate unit authenticates the FortiClient SSL VPN request based on the user group options. The FortiGate unit establishes a tunnel with the client and assigns a virtual IP address to the client PC. Once the tunnel has been established, the user can access the network behind the FortiGate unit.

FortiClient software is available for download at www.forticlient.com and is available for Windows, Mac OS X, Apple iOS, and Android.

Tunnel mode client configuration

The FortiClient SSL VPN tunnel client requires basic configuration by the remote user to connect to the SSL VPN tunnel. When distributing the FortiClient software, provide the following information for the remote user to enter once the client software has been started. Once entered, they can select **Connect** to begin an SSL VPN session.

Connection Name	If you have pre-configured the connection settings, select the connection from the list and then select Connect . Otherwise, enter the settings in the fields below.
Remote Gateway	Enter the IP address or FQDN of the FortiGate unit that hosts the SSL VPN.

Username	Enter your username.
Client Certificate	<p>Use this field if the SSL VPN requires a certificate for authentication.</p> <p>Select the required certificate from the drop-down list. The certificate must be installed in the Internet Explorer certificate store.</p>

The SSL VPN web portal

This chapter explains how to use and configure the web portal features. This chapter is written for end users as well as administrators.

The following topics are included:

Connecting to the FortiGate unit

You can connect to the FortiGate unit using a web browser. The URL of the FortiGate interface may vary from one installation to the next. If required, ask your FortiGate administrator for the URL of the FortiGate unit, and obtain a user name and password. You can connect to the web portal using an Android phone, iPhone, or iPad. The FortiGate unit will display the content of the portal to fit the device's screen.

In addition, if you will be using a personal or group security (X.509) certificate to connect to the FortiGate unit, your web browser may prompt you for the name of the certificate. Your FortiGate administrator can tell you which certificate to select.

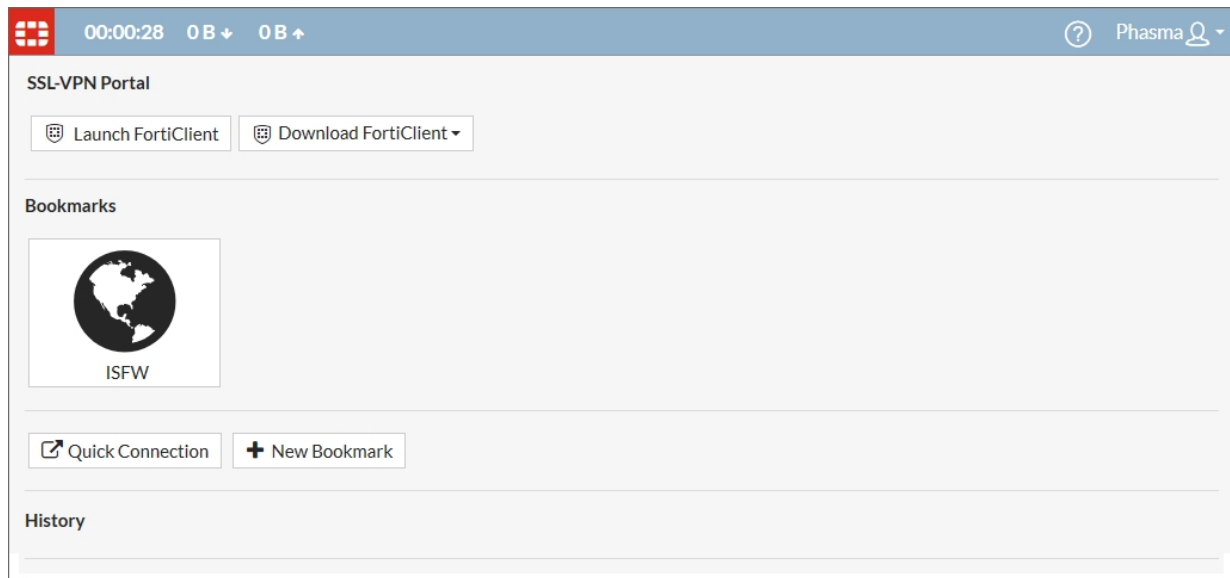
To log into the secure FortiGate HTTP gateway

1. Using the web browser on your computer, browse to the URL of the FortiGate unit (for example, `https://<FortiGate_IP_address>:443/remote/login`). The FortiGate unit may offer you a self-signed security certificate. If you are prompted to proceed, select **Yes**.
A second message may be displayed to inform you that the FortiGate certificate distinguished name differs from the original request. This message is displayed because the FortiGate unit is attempting to redirect your web browser connection. You can ignore the message.
2. When you are prompted for your user name and password:
 - In the **Name** field, type your user name.
 - In the **Password** field, type your password.
3. Select **Login**.
The FortiGate unit will redirect your web browser to the FortiGate SSL VPN web portal home page automatically.

Web portal overview

To log into the web portal, the user must have valid username and password credentials, as well as any other factor of authentication that may be configured, such as a FortiToken code.

After logging in to the web portal, the remote user is presented with a web portal page similar to the following:



Various widgets provide the web portal's features:

- **Session Information** displays the elapsed time since login and the volume of HTTP and HTTPS traffic, both inbound and outbound.
- **Quick Connection** enables you to connect to network resources without using or creating a bookmark.
- **Download Forticlient** provides access to the FortiClient tunnel application for various operating systems.
- **Bookmarks** provides links to network resources. You can use the administrator-defined bookmarks and you can add your own bookmarks.

While using the web portal, you can select the **Help** button to get information to assist you in using the portal features. This information displays in a separate browser window.

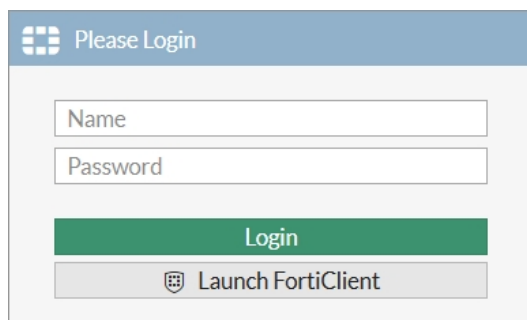
When you have finished using the web portal, select the **Logout** button in the top right corner of the portal window.



After making any changes to the web portal configuration, be sure to select **Apply**.

The Launch FortiClient button

A **Launch FortiClient** button appears on the SSL VPN login page on non-Linux environments when SSL VPN is in tunnel mode.



A similar button is also available within the web portal.

When the user clicks the Launch FortiClient button:

- If FortiClient is installed, FortiClient opens and the GUI switches to the **Remote Access** VPN tab (but FortiClient does *not* automatically create a VPN connection based on the web-mode connection information).
- If FortiClient is not installed, nothing happens.



The **Launch FortiClient** button only works for Mac and Windows operating systems. This feature is not applicable to operating systems that do not support FortiClient, such as Linux and Android phone.

For more information, refer to [The SSL VPN client](#) section.

Portal configuration

The SSL VPN web portal enables users to access network resources through a secure channel using a web browser. Fortinet administrators can configure log in privileges for system users and which network resources are available to the users.

The portal configuration determines what the user sees when they log in to the portal. Both the system administrator and the user have the ability to customize the SSL VPN portal.

There are three pre-defined default web portal configurations available:

- **full-access**: Includes all widgets available to the user - **Session Information**, **Tunnel Mode** options, **Connection Launcher**, **Remote Desktop**, and **Predefined Bookmarks**.
- **tunnel-access**: Includes **Session Information** and **Tunnel Mode** options.
- **web-access**: Includes **Session Information** and **Predefined Bookmarks** widgets.

You can also create your own web portal to meet your corporate requirements.

Portal page	
Create New	Creates a new web portal.
Edit	Select a portal from the list to enable the Edit option, and modify the portal configuration.
Delete	Removes a portal configuration. To remove multiple portals from the list, select the check box beside the portal names, then select Delete .
Name	The name of the web portal.

Portal page

Ref.

Displays the number of times the object is referenced in other configurations on the FortiGate unit, such as security policies.

To view the location of the referenced object, select the number in **Ref.** column.

To view more information about how the object is used, select one of:

View the list page for these objects – automatically redirects you to the list page where the object is referenced at.

Edit this object – modifies settings within that particular setting that the object is referenced with.

View the details for this object – similar to the log viewer table, contains information about what settings are configured within that particular setting that the object is referenced with.

Portal settings

A web portal defines SSL VPN user access to network resources. The portal configuration determines what SSL VPN users see when they log in to the unit. Both the Fortinet administrator and the SSL VPN user have the ability to customize the web portal settings. Portal settings are configured in **VPN > SSL-VPN Portals**.

The following settings are available, allow you to configure general and security console options for your web portal.

Portal Setting	Description
Name	The name for the portal.
Limit Users to One SSL-VPN Connection at a Time	You can set the SSL VPN tunnel such that each user can only log into the tunnel one time concurrently per user per login. That is, once logged into the portal, they cannot go to another system and log in with the same credentials again. This option is disabled by default.
Tunnel Mode	These settings determine how tunnel mode clients are assigned IPv4 addresses.

Portal Setting	Description
Enable Split Tunneling	<p>Select so that the VPN carries only the traffic for the networks behind the FortiGate unit. The user's other traffic follows its normal route.</p> <p>If you enable split tunneling, you are required to set the Routing Address, which is the address that your corporate network is using. Traffic intended for the Routing Address will not be split from the tunnel.</p>
Source IP Pools	Select an IP Pool for users to acquire an IP address when connecting to the portal. There is always a default pool available if you do not create your own.
Tunnel Mode Client Options	<p>These options affect how the FortiClient application behaves when connected to the FortiGate VPN tunnel. When enabled, a check box for the corresponding option appears on the VPN login screen in FortiClient, and is not enabled by default.</p> <ul style="list-style-type: none"> • Allow client to save password - When enabled, if the user selects this option, their password is stored on the user's computer and will automatically populate each time they connect to the VPN. • Allow client to connect automatically - When enabled, if the user selects this option, when the FortiClient application is launched, for example after a reboot or system startup, FortiClient will automatically attempt to connect to the VPN tunnel. • Allow client to keep connections alive - When enabled, if the user selects this option, the FortiClient should try to reconnect once it detects the VPN connection is down unexpectedly (not manually disconnected by user).
Enable Web Mode	Select to enable web mode access.
Portal Message	This is a text header that appears on the top of the web portal.
Theme	Select a color styling specifically for the web portal.
Show Session Information	The Show Session Information widget displays the login name of the user, the amount of time the user has been logged in and the inbound and outbound traffic statistics.
Show Connection Launcher	Displays the Connection Launcher widget in the web portal.
Show Login History	Select to include user login history on the web portal.

Portal Setting	Description
User Bookmarks	Enable to allow users to add their own bookmarks in the web portal.
Predefined Bookmarks	Select to include bookmarks on the web portal. Bookmarks are used as links to internal network resources. When a bookmark is selected from a bookmark list, a pop-up window appears with the web page. Telnet, VNC, and RDP require a browser plugin. FTP and Samba replace the bookmarks page with an HTML file-browser.

Predefined Bookmarks

Bookmarks are used as links to specific resources on the network. When a bookmark is selected from a bookmark list, a pop-up window appears with the requested web page. Telnet, RDP, and VNC pop up a window that requires a browser plug-in. FTP and Samba replace the bookmarks page with an HTML file-browser.

Note that the RDP/VNC web portals are **not** supported for the following platforms:

Platform	Model
FortiGate	80D, 92D, 200D, 200D-POE, 240D, 240D-POE, 600C, 800C, 1000C, 3240C, 3600C, and 5001C
FortiGate-Rugged	90D
FortiWiFi	92D

A web bookmark can include login credentials to automatically log the SSL VPN user into the web site. When the administrator configures bookmarks, the web site credentials must be the same as the user's SSL VPN credentials. Users configuring their own bookmarks can specify alternative credentials for the web site.

Applications available in the web portal

Depending on the web portal configuration and user group settings, one or more of the following server applications are available to you through **Predefined Bookmarks**, as well as the **Quick Connection** widget:

- Citrix makes use of SOCKS so that the Citrix client can connect to the SSL VPN port forward module to provide the connection.
- FTP (File Transfer Protocol) enables you to transfer files between your computer and a remote host.
- HTTP/HTTPS accesses web pages.
- Port Forward provides the middle ground between web mode and tunnel mode. When the SSL VPN receives data from a client application, the data is encrypted and sent to the FortiGate unit, which then forwards the traffic to the application server.
- RDP (Remote Desktop Protocol), similar to VNC, enables you to remotely control a computer running Microsoft Terminal Services.
- SMB/CIFS implements the Server Message Block (SMB) protocol to support file sharing between your computer and a remote server host.
- SSH (Secure Shell) enables you to exchange data between two computers using a secure channel.

- TELNET (Teletype Network emulation) enables you to use your computer as a virtual text-only terminal to log in to a remote host.
- VNC (Virtual Network Computing) enables you to remotely control another computer, for example, accessing your work computer from your home computer.

Some server applications may prompt you for a user name and password. You must have a user account created by the server administrator so that you can log in.



Windows file sharing through SMB/CIFS is supported through shared directories.

Group-based SSL VPN bookmarks

The administrator can add bookmarks for groups of users. SSL VPN will only output the matched group-name entry to the client. This can only be done via the CLI.

To add group-based SSL VPN bookmarks - CLI:

```
config vpn ssl web portal
  edit "portal-name"
    set user-group-bookmark enable*/disable
  next
end
config vpn ssl web user-group-bookmark
  edit "group-name"
    config bookmark
      edit "bookmark1"
        ....
      next
    end
  next
end
```

Downloading files from an SMB server in Web Mode

When logging into the SSL VPN in Web Mode, the client can connect to their file server via SMB and download multiple files at the same time. The client can select/deselect individual files for download, opt to download all, or select all checkboxes for download.

Split DNS support for SSL VPN portals

Split DNS for SSL VPN portals allows you to specify which domains are resolved by the DNS server specified by the VPN, while all other domains are resolved by the DNS specified locally. This feature is useful in both Enterprise and MSP scenarios (when hosting multiple SSL VPN portals).

FortiClient receives this information when the client connects in tunnel mode. FortiClient will push the DNS servers specified to the clients computer and all DNS requests will first attempt use this DNS server. The FortiClient network driver will intercept DNS requests; if they match the `split-dns-domains` listed, the DNS request will go across the tunnel and be resolved by the specified DNS servers.

If the domain does not match `split-dns-domains` then the FortiClient network driver will respond to the DNS request with "no such name" forcing the DNS request to be resolved by the physical adapter DNS.

To configure split DNS support for SSLVPN portals - CLI:

```

config vpn ssl web portal
  edit <name>
    config split-dns-domains
      edit <name>
        set domains "abc.com, cde.com"
        set dns-server1 192.168.1.1
        set dns-server2 192.168.1.2
        set ipv6-dns-server1 xxxxxxxxxxxx
        set ipv6-dns-server2 xxxxxxxxxxxx
      next
    ...
  end
end

```

Using the Bookmarks widget

The Bookmarks widget shows both administrator-configured and user-configured bookmarks. Administrator bookmarks cannot be altered but you can add, edit or delete user bookmarks.

The FortiGate unit forwards client requests to servers on the Internet or internal network. To use the web-portal applications, you add the URL, IP address, or name of the server application to the My Bookmarks list. For more information, see [Adding bookmarks on page 2692](#).



If you want to access a web server or telnet server without first adding a bookmark to the My Bookmarks list, use the Connection Tool instead. For more information, see [Using the Bookmarks widget on page 2692](#).

Adding bookmarks

You can add frequently used connections as bookmarks. Afterward, select any hyperlink from the Bookmarks list to initiate a session.

To add a bookmark

1. In the web portal, select **New Bookmark**.
2. Enter the following information:

Name	Enter the name to display in the Bookmarks list.
Type	Select the abbreviated name of the server application or network service from the drop-down list.
Location	Enter the IP address or FQDN of the server application or network service. For RDP connections, you can append some parameters to control screen size and keyboard layout. See Using the Bookmarks widget on page 2692 .
Description	Optionally enter a short description. The description displays when you pause the mouse pointer over the hyperlink.

SSO	<p>Single Sign On (SSO) is available for HTTP/HTTPS bookmarks only.</p> <p>Disabled — This is not an SSO bookmark.</p> <p>Automatic — Use your SSL VPN credentials or an alternate set. See the SSO Credentials field.</p> <p>Static — Supply credentials and other required information (such as an account number) to a web site that uses an HTML form for authentication. You provide a list of the form field names and the values to enter into them. This method does not work for sites that use HTTP authentication, in which the browser opens a pop-up dialog box requesting credentials.</p>
SSO fields	
SSO Credentials	<p>SSL VPN Login — Use your SSL VPN login credentials.</p> <p>Alternative — Enter Username and Password below.</p>
Username	Alternative username. Available if SSO Credentials is Alternative .
Password	Alternative password. Available if SSO Credentials is Alternative .
Static SSO fields	These fields are available if SSO is Static .
Field Name	Enter the field name, as it appears in the HTML form.
Value	<p>Enter the field value.</p> <p>To use the values from SSO Credentials, enter %passwd% for password or %username% for username.</p>
Add	Add another Field Name / Value pair.

3. Select **OK** and then select **Done**.

Group-based SSL VPN bookmarks

This CLI-only feature allows administrators to add bookmarks for groups of users. SSL VPN will only output the matched group-name entry to the client.

Syntax:

```
config vpn ssl web portal
    edit "portal-name"
        set user-group-bookmark enable*/disable
    next
end
conf vpn ssl web user-group-bookmark
    edit "group-name"
        conf bookmark
            edit "bookmark1"
                ....
            next
        end
```

```
    next
end
```

Group-based SSL VPN bookmarks

This CLI-only feature allows administrators to add bookmarks for groups of users. SSL VPN will only output the matched group-name entry to the client.

Syntax:

```
config vpn ssl web portal
    edit <portal-name>
        set user-group-bookmark [enable | disable]
    next
end
config vpn ssl web user-group-bookmark
    edit <group-name>
        config bookmark
            edit <bookmark1>
                ....
            next
        end
    next
end
```

Automatic bookmarks for SSO credentials

The following CLI changes SSL VPN to send the basic authorization to the remote server for automatic SSO bookmark every time, but only if it is for the same host name. If this attribute is disabled, the SSO credentials are sent to the remote server for every HTTP request.

Syntax

```
config vpn ssl web user-bookmark
    edit <name>
        config bookmarks
            edit <name>
                set sso-credential-sent-once {enable | disable}
            next
        end
    next
end
```

Using the Quick Connection Tool

The **Quick Connection Tool** widget enables a user to connect to a resource when it isn't a predefined bookmark.

You can connect to any type of server without adding a bookmark to the **Bookmarks** list. The fields in the **Quick Connection Tool** enable you to specify the type of server and the URL or IP address of the host computer.

See the following procedures:

- [To connect to a web server on page 2695](#)
- [To ping a host or server behind the FortiGate unit on page 2695](#)

- To start a Telnet session on page 2695
- To start an FTP session on page 2696
- To start an SMB/CIFS session on page 2696
- To start an SSH session on page 2696
- To start an RDP session on page 2697
- To start a VNC session on page 2698

Except for ping, these services require that you have an account on the server to which you connect.



When you use **Quick Connection Tool**, the FortiGate unit may offer you its self-signed security certificate. Select **Yes** to proceed. A second message may be displayed to inform you of a host name mismatch. This message is displayed because the FortiGate unit is attempting to redirect your web browser connection. Select **Yes** to proceed.

To connect to a web server

1. In **Type**, select **HTTP/HTTPS**.
2. In the **Host** field, type the URL of the web server.
For example: `http://www.mywebexample.com` or `https://172.20.120.101`
3. Select **Launch**.
4. To end the session, close the browser window.

To ping a host or server behind the FortiGate unit

1. In **Type**, select **Ping**.
2. In the **Host** field, enter the IP address of the host or server that you want to reach.
For example: `10.11.101.22`
3. Select **Launch**.
A message stating whether the IP address can be reached or not is displayed.

To start a Telnet session

1. In **Type**, select **Telnet**.
2. In the **Host** field, type the IP address of the telnet host.
For example: `10.11.101.12`
3. Select **Launch**.
A Telnet window opens.
4. Select **Connect**.
5. A telnet session starts and you are prompted to log in to the remote host.
After you log in, you may enter any series of valid telnet commands at the system prompt.
6. To end the session, select **Disconnect** (or type `exit`) and then close the TELNET connection window.

To start an FTP session

1. In **Type**, select **FTP**.
2. In the **Host** field, type the IP address of the FTP server.
For example: 10.11.101.12
3. Select **Launch**.
A login window opens.
4. Enter your user name and password and then select **Login**.
You must have a user account on the remote host to log in.
5. Manipulate the files in any of the following ways:
 - To download a file, select the file link in the **Name** column.
 - To access a subdirectory (**Type** is **Folder**), select the link in the **Name** column.
 - To create a subdirectory in the current directory, select **New directory**.
 - To delete a file or subdirectory from the current directory, select its **Delete** icon.
 - To rename a file in the current directory, select its **Rename** icon.
 - To upload a file to the current directory from your client computer, select **Upload**.
 - When the current directory is a subdirectory, you can select **Up** to access the parent directory.
6. To end the FTP session, select **Logout**.

To start an SMB/CIFS session

1. In **Type**, select **SMB/CIFS**.
2. In the **Host** field, type the IP address of the SMB or CIFS server.
For example: 10.11.101.12
3. Select **Launch**.
4. Enter your user name and password and then select **Login**.
You must have a user account on the remote host to log in.
5. Manipulate the files in any of the following ways:
 - To download a file, select the file link in the **Name** column.
 - To access a subdirectory (**Type** is **Folder**), select the file link in the **Name** column.
 - To create a subdirectory in the current directory, select **New Directory**.
 - To delete a file or subdirectory from the current directory, select its **Delete** icon.
 - To rename a file, select its **Rename** icon.
 - To upload a file from your client computer to the current directory, select **Upload**.
 - When the current directory is a subdirectory, you can select **Up** to access the parent directory.
6. To end the SMB/CIFS session, select **Logout** and then close the SMB/CIFS window.

To start an SSH session

1. In **Type**, select **SSH**.
2. In the **Host** field, type the IP address of the SSH host.
For example: 10.11.101.12
3. Select **Launch**.
A login window opens.

4. Select **Connect**.

A SSH session starts and you are prompted to log in to the remote host. You must have a user account to log in. After you log in, you may enter any series of valid commands at the system prompt.

5. To end the session, select **Disconnect** (or type `exit`) and then close the SSH connection window.

To start an RDP session

1. In **Type**, select **RDP**.

2. In the **Host** field, type the IP address of the RDP host.

For example: `10.11.101.12`

3. Optionally, you can specify additional options for RDP by adding them to the **Host** field following the host address. See [RDP options on page 2697](#) for information about the available options.

For example, to use a French language keyboard layout you would add the `-m` parameter:

```
10.11.101.12 -m fr
```

4. To log in to the remote host, type your user name and password. You must have a user account on the remote host to log in. Note that the user name should be entered in User Principal Name (UPN) format.

5. Select **Launch**.

A login window opens.

6. When you see a screen configuration dialog, click **OK**.

The screen configuration dialog does not appear if you specified the screen resolution with the host address.

7. Select **Login**.

If you need to send Ctrl-Alt-Delete in your session, use Ctrl-Alt-End.

8. To end the RDP session, Log out of Windows or select **Cancel** from the Logon window.



Some Windows servers require a specific Security to be set for RDP sessions, such as Network Level Authentication (NLA) or Transport Layer Security (TLS), not the standard RDP encryption security. For example, Windows 10 requires the use of TLS.

RDP options

Locale/Keyboard	-m <locale>			
Use this option if the remote computer might not use the same keyboard layout as your computer. Select the locale code that matches your computer.	The supported values of <locale> are:			
	ar da de de-ch en-gb en-uk en-us es fi fr fr-be fr-ca fr-ch hr hu	Arabic Danish German Swiss German British English UK English US English Spanish Finnish French Belgian French Canadian French Swiss French Croatian Hungarian	it ja lt lv mk no pl pt pt-br ru sl sv tk tr	Italian Japanese Lithuanian Latvian Macedonian Norwegian Polish Portuguese Brazilian Portuguese Russian Slovenian Sudanese Turkmen Turkish

To start a VNC session

1. In **Type**, select **VNC**.
2. In the **Host** field, type the IP address of the VNC host.
For example: 10.11.101.12
3. Select **Launch**.
A login window opens.
4. Type your user name and password when prompted to log in to the remote host.
You must have a user account on the remote host to log in.
5. Select **OK**.
If you need to send Ctrl-Alt-Delete in your session, press F8, then select **Send Ctrl-Alt-Delete** from the pop-up menu.
6. To end the VNC session, close the VNC window.

Note that the RDP/VNC web portals are **not** supported for the following platforms:

Platform	Model
FortiGate	80D, 92D, 200D, 200D-POE, 240D, 240D-POE, 600C, 800C, 1000C, 3240C, 3600C, and 5001C
FortiGate-Rugged	90D
FortiWiFi	92D

Using FortiClient

Remote users can use FortiClient Endpoint Security to initiate an SSL VPN tunnel to connect to the internal network. FortiClient uses local port TCP 1024 to initiate an SSL encrypted connection to the FortiGate unit, on port TCP 10443. When connecting using FortiClient, the FortiGate unit authenticates the FortiClient SSL VPN request based on the user group options. the FortiGate unit establishes a tunnel with the client and assigns a virtual IP address to the client PC. Once the tunnel has been established, the user can access the network behind the FortiGate unit.

For information on configuring the FortiGate unit for SSL VPN connectivity, see [Basic configuration on page 2651](#). For details on configuring FortiClient for SSL VPN connections, see the FortiClient documentation.


Setup examples

The examples in this chapter demonstrate the basic configurations needed for common connections to the SSL VPN tunnel and portals, applying the steps outlined in [Basic configuration on page 2651](#).

The following examples are included:

Secure Internet browsing

This example sets up an SSL VPN tunnel that provides remote users the ability to access the Internet while traveling, and ensures that they are not subject to malware and other dangers, by using the corporate firewall to filter all of their Internet traffic. Essentially, the remote user will connect to the corporate FortiGate unit to surf the Internet.



Using SSL VPN and FortiClient SSL VPN software, you create a means to use the corporate FortiGate to browse the Internet safely.

Creating an SSL VPN IP pool and SSL VPN web portal

1. Go to **VPN > SSL-VPN Portals** and select *tunnel-access*.
2. Disable **Split Tunneling**.
3. For **Source IP Pools** select **SSLVPN_TUNNEL_ADDR1**.
4. Select **OK**.

Creating the SSL VPN user and user group

1. Create the SSL VPN user and add the user to a user group configured for SSL VPN use.
2. Go to **User & Device > User Definition** and select **Create New** to add the user:

User Name	twhite
Password	password

3. Select **OK**.
4. Go to **User & Device > User Groups** and select **Create New** to add *twhite* to a group called **SSL VPN**:

Name	SSL VPN
Type	Firewall

5. Move **twwhite** to the **Members** list.
6. Select **OK**.

Creating a static route for the remote SSL VPN user

Create a static route to direct traffic destined for tunnel users to the SSL VPN tunnel.

1. Go to **Network > Static Routes** and select **Create New** to add the static route.

Destination IP/Mask	10.212.134.0/255.255.255.0
Device	ssl.root



The **Destination IP/Mask** matches the network address of the remote SSL VPN user.

2. Select **OK**.

Creating security policies

Create an SSL VPN security policy with SSL VPN user authentication to allow SSL VPN traffic to enter the FortiGate unit.

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Add an SSL VPN security policy as below, and click **OK**.

Incoming Interface	ssl.root
Outgoing Interface	internal
Source Address	all
Source User Group	SSL VPN
Destination	all

3. Select **OK**.

Configuring authentication rules

1. Go to **VPN > SSL-VPN Settings** and select **Create New** under **Authentication/Portal Mapping**.
2. Add an authentication rule for the remote user:

Users/Groups	Tunnel
Portal	tunnel-access

3. Select **OK** and **Apply**.

Results

Using the FortiClient SSLVPN application, access the VPN using the address `https://172.20.120.136:443/` and log in as `twhite`. Once connected, you can browse the Internet.

From the FortiGate web-based manager, go to **Monitor > SSL-VPN Monitor** to view the list of users connected using SSL VPN. The **Subsession** entry indicates the split tunnel which redirects to the Internet.

Split tunnel

In this configuration, remote users are able to securely access the head office internal network through the head office firewall, yet browse the Internet without going through the head office FortiGate. Split tunneling is enabled by default for SSL VPN on FortiGate units.

The solution below describes how to configure FortiGate SSL VPN split tunneling using the FortiClient SSL VPN software, available from the [Fortinet Support site](#).

Without split tunneling, all communication from remote SSL VPN users to the head office internal network and to the Internet uses an SSL VPN tunnel between the user's PC and the head office FortiGate unit. Connections to the Internet are routed back out the head office FortiGate unit to the Internet. Replies come back into the head office FortiGate unit before being routed back through the SSL VPN tunnel to the remote user.

In short, enabling split tunneling protects the head office from potentially harmful access and external threats that may occur as a result of the end user's indiscretion while browsing the Internet. By contrast, disabling split tunneling protects the end user by forcing all their Internet traffic to pass through the FortiGate firewall.

Creating a firewall address for the head office server

1. Go to **Policy & Objects > Addresses** and select **Create New** and add the head office server address:

Category	Address
Name	Head office server
Type	Subnet
Subnet / IP Range	192.168.1.12
Interface	Internal

2. Select **OK**.

Creating an SSL VPN IP pool and SSL VPN web portal

1. Go to **VPN > SSL-VPN Portals** and select **tunnel-access**.
2. Enter the following:

Name	Connect to head office server
Enable Tunnel Mode	Enable

Enable Split Tunneling	Enable
Routing Address	Internal
Source IP Pools	SSLVPN_TUNNEL_ADDR1

3. Select **OK**.

Creating the SSL VPN user and user group

Create the SSL VPN user and add the user to a user group.

1. Go to **User & Device > User Definition**, select **Create New** and add the user:

User Name	twhite
Password	password

2. Select **OK**.
3. Go to **User & Device > User Groups** and select **Create New** to add the new user to the SSL VPN user group:

Name	Tunnel
Type	Firewall

4. Move **twhite** to the **Members** list.
5. Select **OK**.

Creating a static route for the remote SSL VPN user

Create a static route to direct traffic destined for tunnel users to the SSL VPN tunnel.

1. Go to **Network > Static Routes** and select **Create New**

Destination IP/Mask	10.212.134.0/255.255.255.0
Device	ssl.root

2. Select **OK**.

Creating security policies

Create an SSL VPN security policy with SSL VPN user authentication to allow SSL VPN traffic to enter the FortiGate unit. Create a normal security policy from ssl.root to wan1 to allow SSL VPN traffic to connect to the Internet.

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Complete the following:

Incoming Interface	ssl.root
---------------------------	----------

Source Address	all
Source User(s)	Tunnel
Outgoing Interface	internal
Destination Address	Head office server

3. Select **OK**.
4. Add a security policy that allows remote SSL VPN users to connect to the Internet.
5. Select **Create New**.
6. Complete the following and select **OK**:

Incoming Interface	ssl.root
Source Address	all
Source User(s)	Tunnel
Outgoing Interface	wan1
Destination Address	all
Schedule	always
Service	ALL
Action	ACCEPT

Configuring authentication rules

1. Go to **VPN > SSL-VPN Settings** and select **Create New** under **Authentication/Portal Mapping**.
2. Add an authentication rule for the remote user:

Users/Groups	Tunnel
Portal	tunnel-access

3. Select **OK** and **Apply**.

Results

Using the FortiClient SSL VPN application on the remote PC, connect to the VPN using the address `https://172.20.120.136:443/` and log in with the `twhite` user account. Once connected, you can connect to the head office server or browse to web sites on the Internet.

From the web-based manager, go to **Monitor > SSL-VPN Monitor** to view the list of users connected using SSL VPN. The **Subsession** entry indicates the split tunnel which redirects SSL VPN sessions to the Internet.

Multiple user groups with different access permissions

You might need to provide access to several user groups with different access permissions. Consider the following example topology in which users on the Internet have controlled access to servers and workstations on private networks behind a FortiGate unit.

In this example configuration, there are two users:

- User1 can access the servers on Subnet_1.
- User2 can access the workstation PCs on Subnet_2.

You could easily add more users to either user group to provide them access to the user group's assigned web portal.

General configuration steps

1. Create firewall addresses for:
 - The destination networks.
 - Two non-overlapping tunnel IP address ranges that the FortiGate unit will assign to tunnel clients in the two user groups.
2. Create two web portals.
3. Create two user accounts, User1 and User2.
4. Create two user groups. For each group, add a user as a member and select a web portal. In this example, User1 will belong to Group1, which will be assigned to Portal1 (similar configuration for User2).
5. Create security policies:
 - Two SSL VPN security policies, one to each destination.
 - Two tunnel-mode policies to allow each group of users to reach its permitted destination network.
6. Create the static route to direct packets for the users to the tunnel.

Creating the firewall addresses

Security policies do not accept direct entry of IP addresses and address ranges. You must define firewall addresses in advance.

Creating the destination addresses

SSL VPN users in this example can access either Subnet_1 or Subnet_2.

To define destination addresses - web-based manager:

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**, enter the following information, and select **OK**:

Name	Subnet_1
Type	Subnet
Subnet/IP Range	10.11.101.0/24
Interface	port2

3. Select **Create New**, enter the following information, and select **OK**:

Name	Subnet_2
Type	Subnet
Subnet/IP Range	10.11.201.0/24
Interface	port3

Creating the tunnel client range addresses

To accommodate the two groups of users, split an otherwise unused subnet into two ranges. The tunnel client addresses must not conflict with each other or with other addresses.

To define tunnel client addresses - web-based manager:

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New**, enter the following information, and select **OK**:

Name	Tunnel_group1
Type	IP Range
Subnet/IP Range	10.11.254.1-10.11.254.50
Interface	Any

3. Select **Create New**, enter the following information, and select **OK**.

Name	Tunnel_group2
Type	IP Range
Subnet/IP Range	10.11.254.51-10.11.254.100
Interface	Any

Creating the web portals

To accommodate two different sets of access permissions, you need to create two web portals, portal1 and portal2, for example. Later, you will create two SSL VPN user groups, one to assign to portal1 and the other to assign to portal2.

To create the portal1 web portal:

1. Go to **VPN > SSL-VPN Portals** and select **Create New**.
2. Enter `portal1` in the **Name** field.
3. In **Source IP Pools**, select **Tunnel_group1**.
4. Select **OK**.

To create the portal2 web portal:

1. Go to **VPN > SSL-VPN Portals** and select **Create New**.
2. Enter `portal2` in the **Name** field and select **OK**.
3. In **IP Pools**, select **Tunnel_group2**
4. Select **OK**.

Later, you can configure these portals with bookmarks and enable connection tool capabilities for the convenience of your users.

Creating the user accounts and user groups

After enabling SSL VPN and creating the web portals that you need, you need to create the user accounts and then the user groups that require SSL VPN access.

Go to **User & Device > User Definition** and create user1 and user2 with password authentication. After you create the users, create the SSL VPN user groups.

To create the user groups - web-based manager:

1. Go to **User & Device > User Groups**.
2. Select **Create New** and enter the following information:

Name	Group1
Type	Firewall

3. From the **Available** list, select **User1** and move it to the **Members** list by selecting the right arrow button.
4. Select **OK**.
5. Repeat steps 2 through 4 to create Group2, assigned to Portal2, with User2 as its only member.

Creating the security policies

You need to define security policies to permit your SSL VPN clients, web-mode or tunnel-mode, to connect to the protected networks behind the FortiGate unit. Before you create the security policies, you must define the source and destination addresses to include in the policy. See [Creating the firewall addresses on page 2704](#).

Two types of security policy are required:

- An SSL VPN policy enables clients to authenticate and permits a web-mode connection to the destination network. In this example, there are two destination networks, so there will be two SSL VPN policies. The authentication ensures that only authorized users can access the destination network.
- A tunnel-mode policy is a regular ACCEPT security policy that enables traffic to flow between the SSL VPN tunnel interface and the protected network. Tunnel-mode policies are required if you want to provide tunnel-mode connections for your clients. In this example, there are two destination networks, so there will be two tunnel-mode policies.

To create the SSL VPN security policies - web-based manager:

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information and click **OK**:

Incoming Interface	ssl.root (sslvpn tunnel interface)
Source Address	All
Source User(s)	Group1
Outgoing Interface	port2
Destination Address	Subnet_1
Service	All

3. Select **Create New**.
4. Enter the following information:

Incoming Interface	ssl.root (sslvpn tunnel interface)
Source Address	All
Source User(s)	Group2
Outgoing Interface	port3
Destination Address	Subnet_2
Service	All

5. Click **OK**.

Configuring authentication rules

1. Go to **VPN > SSL-VPN Settings** and select **Create New** under **Authentication/Portal Mapping**.
2. Add an authentication rule for the first remote group:

Users/Groups	Group1
Portal	Portal1

3. Select **OK** and **Apply**.
4. Select **Create New** and add an authentication rule for the second remote group:

Users/Groups	Group2
Portal	Portal2

5. Select **OK** and **Apply**.

To create the tunnel-mode security policies - web-based manager:

1. Go to **Policy & Objects > IPv4 Policy** and select **Create New**.
2. Enter the following information, and select **OK**:

Incoming Interface	ssl.root (sslvpn tunnel interface)
Source Address	Tunnel_group1
Source User(s)	Group1
Outgoing Interface	port2
Destination Address	Subnet_1
Service	All
Action	ACCEPT
Enable NAT	Enable

3. Select **Create New**.
4. Enter the following information, and select **OK**:

Incoming Interface	ssl.root (sslvpn tunnel interface)
Source Address	Tunnel_group2
Source User(s)	Group2
Outgoing Interface	port3
Destination Address	Subnet_2
Service	All
Action	ACCEPT
Enable NAT	Enable

Create the static route to tunnel mode clients

Reply packets destined for tunnel mode clients must pass through the SSL VPN tunnel. You need to define a static route to allow this.

To add a route to SSL VPN tunnel mode clients - web-based manager:

1. Go to **Network > Static Routes** and select **Create New**.
2. Enter the following information and select **OK**.

Destination IP/Mask	10.11.254.0/24 This IP address range covers both ranges that you assigned to SSL VPN tunnel-mode users. See Creating the tunnel client range addresses on page 2705 .
Device	Select the SSL VPN virtual interface, ssl.root for example.



In this example, the **IP Pools** field on the **VPN > SSL-VPN Settings** page is not used because each web portal specifies its own tunnel IP address range.

Client device certificate authentication with multiple groups

In the following example, we require clients connecting to a FortiGate SSL VPN to have a device certificate installed on their machine in order to authenticate to the VPN.

Employees (in a specific OU in AD) will be required to have a device certificate to connect, while vendors (in a separate OU in AD) will *not* be required to have a device certificate.

This can *only* be performed in the CLI console.



The Authentication-rule option is only available in the CLI as an advanced setting to achieve your requirements. It is not available on the GUI. So in **VPN > SSL-VPN Settings**, do *not* enable **Require Client Certificate**, but selectively enable `client-cert` in each authentication-rule based on the requirements through CLI instead.

Configuring SSL VPN shared settings and authentication rules - CLI:

The following example assumes that remote LDAP users/groups have been pre-configured.

```
config vpn ssl settings
  set servercert "Fortinet_Factory"
  set tunnel-ip-pools "SSLVPN_TUNNEL_ADDR1"
  set port 443
  set source-interface "wan1"
  set source-address "all"
  set default-portal "full-access"
  config authentication-rule
    edit 1
      set source-interface "wan1"
      set source-address "all"
      set groups "Employees"
      set portal "full-access"
      set client-cert enable
    next
    edit 2
      set source-interface "wan1"
      set source-address "all"
      set groups "Vendors"
      set portal "full-access"
      set client-cert disable <-- Set by default and will not be displayed.
    next
  end
end
```

Configure the remainder of the SSL VPN tunnel as normal (creating a firewall policy allowing SSL VPN access to the internal network, including the VPN groups, necessary security profiles, etc.).

If configured correctly, only the 'Employees' group should require a client certificate to authenticate to the VPN.

Troubleshooting

This section contains tips to help you with some common challenges of SSL VPNs, and other miscellaneous information useful for testing and troubleshooting.

- Enter the following to display debug messages for SSL VPN:

```
diagnose debug application sslvpn -1
```

This command enables debugging of SSL VPN with a debug level of -1. The -1 debug level produces detailed results.

- Enter the following command to verify the debug configuration:

```
diagnose debug info
debug output: disable
console timestamp: disable
console no user log message: disable
sslvpn debug level: -1 (0xffffffff)
CLI debug level: 3
```

This output verifies that SSL VPN debugging is enabled with a debug level of -1, and shows what filters are in place. The output above indicates that debug output is disabled, so debug messages are not displayed. The output also indicates that debugging has not been enabled for any software systems.

- Enter the following to enable displaying debug messages:

```
diagnose debug enable
```

To view the debug messages, log into the SSL VPN portal. The CLI displays debug output similar to the following:

```
FGT60C3G10002814 # [282:root]SSL state:before/accept initialization (172.20.120.12)
[282:root]SSL state:SSLv3 read client hello A (172.20.120.12)
[282:root]SSL state:SSLv3 write server hello A (172.20.120.12)
[282:root]SSL state:SSLv3 write change cipher spec A (172.20.120.12)
[282:root]SSL state:SSLv3 write finished B (172.20.120.12)
[282:root]SSL state:SSLv3 flush data (172.20.120.12)
[282:root]SSL state:SSLv3 read finished A:system lib(172.20.120.12)
[282:root]SSL state:SSLv3 read finished A (172.20.120.12)
[282:root]SSL state:SSL negotiation finished successfully (172.20.120.12)
[282:root]SSL established: DHE-RSA-AES256-SHA SSLv3 Kx=DH Au=RSA Enc=AES(256) Mac=SHA1
```

- Enter the following to stop displaying debug messages:

```
diagnose debug disable
```

The following is a list of potential issues. The suggestions below are not exhaustive, and may not reflect your network topology.

There is no response from the SSL VPN URL.

- Go to **VPN > SSL-VPN Settings** and check the SSL VPN port assignment. Also, verify that the SSL VPN policy is configured correctly.
- Check the URL you are attempting to connect to. It should follow this pattern:

```
https://<FortiGate IP>:<Port>/remote/login
```

- Ensure that you are using the correct port number in the URL.

FortiClient cannot connect.

Read the Release Notes to ensure that the version of FortiClient you are using is compatible with your version of FortiOS.

Tunnel-mode connection shuts down after a few seconds.

This issue can occur when there are multiple interfaces connected to the Internet (for example, a dual WAN). Upgrade to the latest firmware then use the following CLI command:

```
config vpn ssl settings
    set route-source-interface enable
end
```

When you attempt to connect using FortiClient or in Web mode, you are returned to the login page, or you receive the following error message: “Unable to logon to the server. Your user name or password may not be configured properly for this connection. (-12).”

- Ensure that cookies are enabled in your browser.
- If you are using a remote authentication server, ensure that the FortiGate is able to communicate with it.
- Access to the web portal or tunnel will fail if Internet Explorer has the privacy Internet Options set to High. If set to High, Internet Explorer will block cookies that do not have a compact privacy policy, and that use personally identifiable information without your explicit consent.

You receive an error message stating: “Destination address of Split Tunneling policy is invalid.”

The SSL VPN security policy uses the **ALL** address as its destination. Change the address to that of the protected network instead.

The tunnel connects but there is no communication.

Go to **Network > Static Routes** and ensure that there is a static route to direct packets destined for the tunnel users to the SSL VPN interface.

You can connect remotely to the VPN tunnel but are unable to access the network resources.

Go to **Policy & Objects > IPv4 Policy** and examine the policy allowing VPN access to the local network. If the destination address is set to all, create a firewall address for the internal network. Change the destination address and attempt to connect remotely again.

Users are unable to download the SSL VPN plugin.

Go to **VPN > SSL-VPN Portals** to make sure that the option to **Limit Users to One SSL-VPN Connection at a Time** is disabled. This allows users to connect to the resources on the portal page while also connecting to the VPN through FortiClient.

Users are being assigned to the wrong IP range.

Ensure that the same IP Pool is used in VPN Portal and VPN Settings to avoid conflicts. If there is a conflict, the portal settings will be used.

Flow-based (vdom) AntiVirus profiles in SSL VPN web mode limitation

In flow mode vdom, SSL VPN web mode doesn't block antivirus even when `av-profile` is set (however, SSL VPN tunnel mode AV profile does work).

Sending tunnel statistics to FortiAnalyzer

By default, logged events include tunnel-up and tunnel-down status events. Other events, by default, will appear in the FortiAnalyzer report as "No Data Available". More accurate results require logs with `action=tunnel-stats`, which is used in generating reports on the FortiAnalyzer (rather than the tunnel-up and tunnel-down event logs). The FortiGate does not, by default, send `tunnel-stats` information.

To allow VPN `tunnel-stats` to be sent to FortiAnalyzer, configure the FortiGate unit as follows using the CLI:

```
config system settings
    set vpn-stats-log ipsec ssl
    set vpn-stats-period 300
end
```

HTTP header information

The X-Content-Type-Options header is added to internal pages of SSL VPN to comport with PCI-DSS compatibility. Strict-Transport-Security is added to the HTTP header for the same reason.

Visibility of SSL VPN portal SSO credentials

The SSL VPN portal SSO feature sends passwords and usernames in clear text to the client in a javascript file. This feature sees the introduction of a new CLI command to control the visibility of these SSO credentials. Enable this command to prevent SSO credentials from being sent to the client.

Syntax

```
config vpn ssl web portal
    edit portal
        set hide-sso-credential {enable | disable}
    next
end
```

Chapter 24 - System Administration

This guide contains the following sections:

- [What's New in FortiOS 6.0](#) informs you about new system administration features in FortiOS 6.0.
- [Administrators](#) describes the tasks that can be done to add and secure administrative access to a FortiGate.
- [Monitoring](#) discusses the various methods of monitoring both your FortiGate and network traffic through a range of different tools available within FortiOS.
- [Replacement messages](#) explains how to view and customize replacement messages on your FortiGate.
- [Administration for schools](#) shares basic practices administrators in school systems can employ .
- [PPTP and L2TP](#) contains information on configuring PPTP and L2TP VPNs as well as PPTP passthrough.
- [Session helpers](#) explains how to use session helpers to analyze data in the packet bodies of some protocols and make adjustments to allow those protocols to send packets through the firewall.
- [Advanced concepts](#) describes more involved administrative topics to enhance network security and traffic efficiency.

What's new in FortiOS 6.0

The following list contains new system administration features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- [Updates to "Administrator profiles" on page 2714](#)
- ["Restrict local admin authentication when remote authentication server is running" on page 2721](#)
- ["FortiToken extension to comply with PCI 3.2" on page 2726](#)
- ["Multi-dimension tagging" on page 2779](#)

Administrators

By default, the FortiGate has a super administrator account, called `admin`, which can't be deleted. Additional administrators can be added for various functions, each with a unique user name, password, and set of access privileges.

The following sections explain how to add and secure administrative access to a FortiGate:

- [Administrator profiles](#)
- [Adding a local administrator](#)
- [LDAP authentication for administrators](#)
- [Other methods of administrator authentication](#)
- [Administrator logout](#)
- [Monitoring administrators](#)
- [Management access](#)
- [Security precautions](#)

Administrator profiles

Administrator profiles define what the administrator can do when logged into the FortiGate. When you set up an administrator account, you also assign an administrator profile dictating what the administrator sees. Depending on the nature of the administrator's work, access level or seniority, you can allow them to view and configure as much, or as little, as required.

super_admin profile

This profile has access to all components of FortiOS, including the ability to add and remove other system administrators. For certain administrative functions, such as backing up and restoring the configuration, `super_admin` access is required. To ensure that there is always a method to administer the FortiGate, the `super_admin` profile can't be deleted or modified.



Lower level administrator profiles can't backup or restore the FortiOS configuration.

The `super_admin` profile is used by the default `admin` account. We recommend that you add a password and rename this account once you have set up your FortiGate. In order to rename the default account, a second admin account is required. For more information, see ["Adding a local administrator" on page 2715](#).

Creating profiles

To configure administrator profiles, go to **System > Admin Profiles** and select **Create New**.

On the **New Admin Profile** page, select the access permissions for the admin profile you are creating. For example, you can configure a profile so that the administrator can only read/write **Firewall** configuration, which includes firewall policies, addresses, services, schedules, packet capture, and some other parts of the FortiGate.

configuration. Any other aspects of the FortiGate configuration, including VPNs and security profiles, would be hidden from this administrator.

Access control can also be set to **Custom** for some features. This allows for more granular control of administrator access. Using the **Firewall** example, you can set access to **Custom** and then select separate **Read/Write** privileges for **Policy**, **Address**, **Service**, and **Schedule**.

Administrator timeout override per access profile

You can configure administrator profiles to increase inactivity timeout and facilitate use of the GUI for central monitoring. This feature allows the `admintimeout` value, under `config system accprofile`, to be overridden per access profile.

Note that you can set this on a per-profile basis, to avoid the option from being unintentionally set globally.

CLI Syntax - Configure admin timeout

```
config system accprofile
edit <name>
    set admintimeout-override {enable | disable}
    set admintimeout <0-480> - (default = 10, 0 = unlimited)
next
end
```

Adding a local administrator

Only administrators with read-write privileges for **User & Device** can create a new administrator account.

To add an administrator - GUI

1. Go to **System > Administrators**.
2. Select **Create New > Administrator**.
3. Add a **Username** for the administrator.



Don't include the characters `<>()#"'` in the administrator's name. Using these characters in the administrator account name can result in a cross site scripting (XSS) vulnerability.

4. Set **Type** to **Local User**.
5. Enter the **Password** for the user. This may be a temporary password that the administrator can change later. Passwords can be up to 256 characters in length. For more information on passwords, see the [Passwords](#) discussion in the Getting Started chapter.
6. Determine if you need to enable security options: **SMS**, **Two-factor Authentication**, **Restrict login to trusted hosts**, **Restrict admin to guest account provisioning only**.
7. Select **OK**.



You can configure guest management administrator's through the GUI. To create the user group to be used for guest user accounts, go to [Managing Guest Access](#) in the [Authentication](#) chapter.

To add an administrator - CLI

```
config system admin
  edit <admin_name>
    set password <password>
    set accprofile <profile_name>
    set guest-auth {enable | disable}
    set user-groups <group-name>
  end
```

The CLI command `set user-groups` can only be used when `guest-auth` is set to `enable`.

LDAP authentication for administrators

Administrators can use remote authentication, such as LDAP, to connect to the FortiGate.

To do this, you must follow these three steps:

- configure the LDAP server
- add the LDAP server to a user group
- configure the administrator account

Configure the LDAP server

First set up the LDAP server as you normally would, and include a group to bind to.

To configure the LDAP server - GUI

1. Go to **User & Device > LDAP Servers** and select **Create New**.
2. Enter a **Name** for the server.
3. Enter the **Server IP** address or name.
4. Enter the **Common Name Identifier** and **Distinguished Name**.
5. Set the **Bind Type** to **Regular** and enter the **Username** and **Password**.
6. Select **OK**.

To configure the LDAP server - CLI

```
config user ldap
  edit <ldap_server_name>
    set server <server_ip>
    set cnid cn
    set dn DC=XYZ,DC=COM
    set type regular
    set user name CN=Administrator,CN=Users,DC=XYZ,DC=COM
    set password <password>
    set member-attr <group_binding>
  end
```

Add the LDAP server to a user group

Next, create a user group that will include the LDAP server that was created above.

To create a user group - GUI

1. Go to **User & Device > User Groups** and select **Create New**.
2. Enter a **Name** for the group.
3. In the section labeled **Remote groups**, select **Create New**.
4. Select the **Remote Server** from the drop-down list.
5. Select **OK**.

To create a user group - CLI

```
config user group
  edit <group_name>
    config match
      edit 1
        set server-name <LDAP_server>
        set group-name <group_name>
      end
    end
  end
```

Configure the administrator account

Now you can create a new administrator, where rather than entering a password, you will use the new user group and the wildcard option for authentication.

To create an administrator - GUI

1. Go to **System > Administrators** and select **Create New**.
2. In the **Administrator** field, enter the name for the administrator.
3. For **Type**, select **Match a user on a remote server group**.
4. Select the **User Group** created above from the drop-down list.
5. Select **Wildcard**. The Wildcard option allows for LDAP users to connect as this administrator.
6. Select an **Admin Profile**.
7. Select **OK**.

To create an administrator - CLI

```
config system admin
  edit <admin_name>
    set remote-auth enable
    set accprofile super_admin
    set wild card enable
    set remote-group ldap
  end
```

Other methods of administrator authentication

Admin accounts can use a variety of methods for authentication, including RADIUS, TACACS+, and PKI.

RADIUS authentication for administrators

If you want to use a RADIUS server to authenticate administrators, you must:

- configure the FortiGate to access the RADIUS server
- create the RADIUS user group
- configure an administrator to authenticate with a RADIUS server.

TACACS+ authentication for administrators

If you want to use a TACACS+ server to authenticate administrators, you must:

- configure the FortiGate to access the TACACS+ server
- create a TACACS+ user group
- configure an administrator to authenticate with a TACACS+ server.

PKI certificate authentication for administrators

To use PKI authentication for an administrator, you must:

- configure a PKI user
- create a PKI user group
- configure an administrator to authenticate with a PKI certificate.

Administrator lockout

By default, the FortiGate sets the number of password retries at three, allowing the administrator a maximum of three attempts to log into their account before locking the account for a set amount of time.

Both the number of attempts (`admin-lockout-threshold`) and the wait time before the administrator can try to enter a password again (`admin-lockout-duration`) can be configured within the CLI.

To configure the lockout options:

```
config system global
    set admin-lockout-threshold <failed_attempts>
    set admin-lockout-duration <seconds>
end
```

The default value of `admin-lockout-threshold` is 3 and the range of values is between 1 and 10. The `admin-lockout-duration` is set to 60 seconds by default and the range of values is between 1 and 4294967295 seconds.

Keep in mind that the higher the lockout threshold, the higher the risk that someone may be able to break into the FortiGate.

Example:

To set the `admin-lockout-threshold` to one attempt and the `admin-lockout-duration` to a five minute duration before the administrator can try to log in again, enter the commands:

```
config system global
    set admin-lockout-threshold 1
    set admin-lockout-duration 300
end
```



If the time span between the first failed login attempt and the `admin-lockout-threshold` failed login attempt is less than `admin-lockout-duration`, the lockout will be triggered.

Monitoring administrators

You can view the administrators logged in using the **System Information** widget on the Dashboard. The **Current Administrator** row that shows the administrator logged in and the total number of administrators logged in. Selecting **Details** displays the administrators, where they are logging in from and how (CLI, GUI) and when they logged in.

You are also able to monitor the activities the administrators perform on the FortiGate using the event logging. Event logs include a number of options to track configuration changes.

To set logging - GUI

1. Go to **Log & Report > Log Settings**.
2. Under **Event Logging**, select **Customize** and ensure **System activity event** is selected.
3. Select **Apply**.

To set logging - CLI

```
config log eventfilter
    set event enable
    set system enable
end
```

To view the logs go to **Log & Report > System Events**.

Management access

Management access defines how administrators are able to log on to the FortiGate. In NAT mode, access is configured for each of the FortiGate's interfaces, using the interface's IP to connect. In transparent mode, a single management IP address is configured to allow access.

Management access can be via HTTP, HTTPS, Telnet, or SSH sessions. HTTPS and SSH are preferred as they are more secure. The management computer must connect to an interface that permits management access and its IP address must be on the same network. If you are using VDOMs, an administrator who is restricted to a specific VDOM must use a computer that connects to an interface on that VDOM.

You can allow remote administration of the FortiGate; however, it is not recommended, since it could compromise the security of the FortiGate. If you require remote administration, the following precautions can be taken to improve the security of a FortiGate:

- Use secure administrator passwords.
- Change these passwords regularly.
- Enable two-factor authentication for administrators.
- Enable secure administrative access to this interface using only HTTPS or SSH.
- Use Trusted Hosts to limit where the remote access can originate from.
- Don't change the system idle timeout from the default value of 5 minutes.

Security precautions

One potential point of a security breach is at the management computer. Administrators who leave their workstations for a prolonged amount of time while staying logged into the GUI or CLI leave the firewall open to malicious intent.



When logging in using a local admin with the default or empty password, a warning prompt will appear upon login. Admins will be logged out if they have no permissions.

Restrict logins from trusted hosts

Setting up trusted hosts for an administrator limits the addresses from where they can log into FortiOS. The trusted hosts configuration applies to most forms of administrative access including HTTPS, SSH, and SNMP. When you identify a trusted host for an administrator account, FortiOS accepts that administrator's login only from one of the trusted hosts. A login, even with proper credentials, from a non-trusted host is dropped.



Even if you have configured trusted hosts, if you have enabled ping administrative access on a FortiGate interface, it will respond to ping requests from any IP address.

To identify trusted hosts, go to **System > Administrators**, edit the administrator account, enable **Restrict login to trusted hosts**, and add up to ten trusted host IP addresses.

To add two trusted hosts from the CLI:

```
config system admin
  edit <administrator-name>
    set trustedhost1 172.25.176.23 255.255.255.255
    set trustedhost2 172.25.177.0 255.255.255.0
  end
```

Trusted host IP addresses can identify individual hosts or subnets. Just like firewall policies, FortiOS searches through the list of trusted hosts in order and acts on the first match it finds. When you configure trusted hosts, start by adding specific addresses at the top of the list. Follow with more general IP addresses. You don't have to add addresses to all of the trusted hosts as long as all specific addresses are above all of the 0.0.0.0 0.0.0.0 addresses.

Prevent concurrent administrator sessions

Concurrent administrator sessions occur when multiple people concurrently access the FortiGate using the same administrator account. This is allowed by default. If you wish to prevent this behavior go to **System > Settings** and disable **Allow multiple concurrent sessions for each administrator**.

From the CLI:

```
config system global
  set admin-concurrent disable
end
```

Note, if you disable concurrent sessions for an administrator, you will be allowed only one session with the same username even if it is from the same IP.

Restrict local admin authentication when remote authentication server is running

The following command can be enabled so that whenever any remote server (TACACS, LDAP, or RADIUS) is up and running, any local admin authentication will be blocked. Local admins will be allowed access only if no remote server is detected.

Syntax:

```
config system global
    set admin-restrict-local {enable | disable} - (Default is set to disable)
end
```

Segregate administrative roles

To minimize the effect of an administrator causing errors to the FortiGate configuration and possibly jeopardizing the network, create individual administrative roles where none of the administrators have super_admin permissions. For example, one account is used solely to create security policies, another for users and groups, another for VPN, and so on.

SSH log in time out

You can take up to 120 seconds to log into the FortiGate when using SSH. You can use the following CLI command to reduce this time and enhance security:

```
config system global
    set admin-ssh-grace-time <number_of_seconds>
end
```

The range can be between 10 and 3600 seconds.

HTTPS redirect

When configuring the Administration Settings (found at **System > Settings**), you can also enable HTTP to **Redirect to HTTPS**. When enabled, if a administrator tries to connect to an interface using HTTP, this traffic will be automatically redirected to use HTTPS instead for a more secure connection.

Administrator log in disclaimers

FortiOS can display a disclaimer before or after logging into the GUI or CLI (or both). In either case the administrator must read and accept the disclaimer before they can proceed.

Use the following command to display a disclaimer before logging in:

```
config system global
    set pre-login-banner enable
end
```

Use the following command to display a disclaimer after logging in:

```
config system global
    set post-login-banner enable
end
```

You can customize the replacement messages for these disclaimers by going to **System > Replacement Messages**. Select **Extended View** to view and edit the **Administrator** replacement messages.

From the CLI:


```
config system replacemsg admin pre_admin-disclaimer-text
config system replacemsg admin post_admin-disclaimer-text
```

Disable the console interface

You can disable your FortiGate's console interface to prevent any unwanted login attempts:

```
config system console
    set login disable
end
```

Disable other interfaces

If any of the interfaces on the FortiGate aren't being used, disable traffic on that interface. This avoids someone plugging in network cables and potentially causing network bypass or loop issues.

To disable an interface - GUI

1. Go to **Network > Interfaces**.
2. Select the interface from the list and select **Edit**.
3. For **Administrative Access**, select **Down**.
4. Select **OK**.

To disable an interface - CLI

```
config system interface
    edit <interface_name>
        set status down
    end
```

Self-signed GUI certificates

For increased security, the self-sign certificate is the default GUI certificate, if the BIOS certificate is using SHA-1.

Monitoring

With network administration, the first step is installing and configuring the FortiGate to be the protector of the internal network. Once the system is running efficiently, the next step is to monitor the system and network traffic. When a threat or vulnerability is discovered, you can make configuration changes as necessary.

This chapter discusses the various methods of monitoring both the FortiGate and the network traffic through a range of different tools available within FortiOS.

This section includes the topics:

- [Dashboard](#)
- [sFlow](#)
- [Monitor menus](#)
- [Logging](#)
- [Alert email](#)
- [SNMP](#)
- [SNMP get command syntax](#)

Dashboard

The FortiOS dashboard provides real-time system information presented in a network operations center (NOC) view with a focus on alerts. By default, the dashboard displays key statistics for the FortiGate, such as memory and CPU status, port health, whether they are up or down and their throughput. Widgets are interactive. By clicking or hovering over most widgets, you can get additional information or follow links to other pages.

The dashboard and its widgets include:

- Multiple dashboard support.
- VDOM and global dashboards.
- Widget resize control.
- Notifications on the top header bar.

For more information, see the [Dashboard](#) discussion in the [Getting Started](#) chapter of the Handbook.

sFlow support

sFlow is a method of monitoring the traffic on your network to identify areas on the network that may impact performance and throughput. FortiOS implements sFlow version 5.

sFlow uses packet sampling to monitor network traffic. The sFlow Agent captures packet information at defined intervals and sends them to an sFlow Collector for analysis, providing real-time data analysis. The information sent is only a sampling of the data for minimal impact on network throughput and performance.

The sFlow Agent is embedded in the FortiGate. Once configured, the FortiGate sends sFlow datagrams of the sampled traffic to the sFlow Collector, also called an sFlow Analyzer. The sFlow Collector receives the datagrams, and provides real-time analysis and graphing to indicate where potential traffic issues are occurring. sFlow Collector software is available from a number of third party software vendors.

sFlow data captures only a sampling of network traffic, not all traffic like the traffic logs on the FortiGate. Sampling works by the sFlow Agent looking at traffic packets when they arrive on an interface. A decision is made

whether the packet is dropped and allowed to be to its destination or if a copy is forwarded to the sFlow Collector. The sample used and its frequency are determined during configuration.

sFlow is not supported on virtual interfaces such as vdom link, ipsec, ssl.root or gre.

The sFlow datagram sent to the Collector contains the information:

- Packet header (e.g. MAC, IPv4, IPv6, IPX, AppleTalk, TCP, UDP, ICMP)
- Sample process parameters (rate, pool, etc.)
- Input/output ports
- Priority (802.1p and TOS)
- VLAN (802.1Q)
- Source/destination prefix
- Next hop address
- Source AS, Source Peer AS
- Destination AS Path
- Communities, local preference
- User IDs (TACACS/RADIUS) for source/destination
- URL associated with source/destination
- Interface statistics (RFC 1573, RFC 2233, and RFC 2358)

sFlow agents can be added to any type of FortiGate interface. sFlow isn't supported on some virtual interfaces such as VDOM link, IPsec, gre, and ssl.root.

For more information on sFlow, Collector software and sFlow MIBs, visit www.sflow.org.

Configuration

sFlow configuration is available only in the CLI. Configuration requires two steps: enabling the sFlow Agent and configuring the interface for the sampling information.

To enable sFlow - CLI:

```
config system sflow
  set collector-ip <ip_address>
  set collector-port <port_number>
  set source-ip <ip_address>
end
```

The default port for sFlow is UDP 6343.

To configure in VDOM - CLI:

```
config system vdom-sflow
  set vdom-sflow enable
  set collector-ip <ip_address>
  set collector-port <port_number>
  set source-ip <ip_address>
end
```

To configure sFlow agents per interface - CLI:

```
config system interface
```

```
edit <interface_name>
    set sflow-sampler enable
    set sample-rate <every_n_packets>
    set sample-direction [tx | rx | both]
    set polling-interval <seconds>
next
end
```

Monitor menus

The **Monitor** menus enable you to view session and policy information and other activity occurring on your FortiGate unit. The monitors provide the details of user activity, traffic and policy usage to show live activity. Monitors are available for DHCP, routing, security policies, traffic shaping, load balancing, security features, VPN, users, and WiFi.

Logging

FortiOS provides a robust logging environment that enables you to monitor, store, and report traffic information and FortiGate events, including attempted log ins and hardware status. Depending on your requirements, you can log to a number of different hosts.

To configure logging in the web-based manager, go to **Log & Report > Log Settings**.

To configure logging in the CLI use the commands `config log <log_location>`.

For details on configuring logging see the [Logging and Reporting Guide](#).

If you will be using several FortiGate units, you can also use a FortiAnalyzer unit for logging. For more information, see the FortiAnalyzer Administration Guide.

Syslog server

An industry standard for collecting log messages, for off-site storage. In the web-based manager, you are able to send logs to a single syslog server, however in the CLI you can configure up to three syslog servers where you can also use multiple configuration options. For example, send traffic logs to one server, antivirus logs to another. The FortiGate unit sends Syslog traffic over UDP port 514. Note that if a secure tunnel is configured for communication to a FortiAnalyzer unit, then Syslog traffic will be sent over an IPsec connection, using UDP 500/4500, Protocol IP/50.

To configure a Syslog server in the web-based manager, go to **Log & Report > Log Settings**. In the CLI use the commands:

```
config log syslogd setting
    set status enable
    set server <IP address or FQDN of syslog server>
end
```

Further options are available when enabled to configure a different port, facility and server IP address.

For Syslog traffic, you can identify a specific port/IP address for logging traffic. Configuration of these services is performed in the CLI, using the command `set source-ip`. When configured, this becomes the dedicated port to send this traffic over.

For example, to set the source IP of a Syslog server to be on the DMZ1 port with an IP of 192.168.4.5, the commands are:

```
config log syslogd setting
    set status enable
    set source-ip 192.168.4.5
end
```

FortiToken extension to comply with PCI 3.2

For FortiToken to be PCI 3.2 compliant, multi-factor authentication with FortiOS can be globally enforced for all login methods under `config system global`.

When `multi-factor-authentication` is set to `mandatory`, the system will collect and log each factor (username, password, and OTP) after authentication.

Note that even if a user is not configured with two-factor authentication, an empty OTP (or any OTP entered) will make second factor authentication pass.

Syntax:

```
config system global
    set multi-factor-authentication {optional | mandatory} - (Default is set to optional)
end
```

Alert email

As an administrator, you want to be certain you can respond quickly to issues occurring on your network or on the FortiGate unit. Alert emails provide an efficient and direct method of notifying an administrator of events. By configuring alert messages, you can define the threshold when a problem becomes critical and needs attention. When this threshold is reached, the FortiGate unit will send an email to one or more individuals, notifying them of the issue.

In the following example, the FortiGate unit is configured to send email to two administrators (admin1 and admin2) when multiple intrusions are detected every two minutes. The FortiGate unit has its own email address on the mail server.

To configure the email service

1. Go to **System > Advanced**.
2. In the **Email Service**, enable **Use Custom Email Server**, complete the following and select **Apply**:

SMTP Server	Enter the address or name of the email server. For example, <code>smtp.example.com</code> .
Default Reply To	Enter an email address to associate with the alert email. This field is optional. If you enter an email address here, it overrides the email address entered when configuring alert email in Log & Report > Email Alert Settings .
Authentication	Enable authentication if required by the email server.
Security mode	Choose between <i>None</i> , <i>SMTPS</i> or <i>STARTTLS</i>
Port	25

To configure alert email - GUI

1. Go to **Log & Report > Email Alert Settings**.
2. Enter the information:

Email from	fortigate@example.com
Email to	admin1@example.com
	admin2@example.com

3. For the **Interval Time**, enter 2.
4. Select **Intrusion Detected**.
5. Select **Apply**.

To configure alert email - CLI

```
config system email-server
  set port 25
  set server smtp.example.com
  set authenticate enable
  set username FortiGate
  set password *****
end
config alertemail setting
  set username fortigate@example.com
  set mailto1 admin1@example.com
  set mailto2 admin2@example.com
  set filter category
  set IPS-logs enable
end
```

SNMP

Simple Network Management Protocol (SNMP) enables you to monitor hardware on your network. You can configure the hardware, such as the FortiGate SNMP agent, to report system information and send traps (alarms or event messages) to SNMP managers. An SNMP manager, or host, is typically a computer running an application that can read the incoming trap and event messages from the agent and send out SNMP queries to the SNMP agents. A FortiManager unit can act as an SNMP manager to one or more FortiGate units. FortiOS supports SNMP using IPv4 and IPv6 addressing.

By using an SNMP manager, you can access SNMP traps and data from any FortiGate interface or VLAN subinterface configured for SNMP management access. Part of configuring an SNMP manager is to list it as a host in a community on the FortiGate unit it will be monitoring. Otherwise, the SNMP monitor will not receive any traps from that FortiGate unit or be able to query that unit.

The FortiGate SNMP implementation is read-only. SNMP v1, v2c, and v3 compliant SNMP managers have read-only access to FortiGate system information through queries and can receive trap messages from the FortiGate unit.

To monitor FortiGate system information and receive FortiGate traps, you must first compile the Fortinet and FortiGate Management Information Base (MIB) files. A MIB is a text file that describes a list of SNMP data

objects that are used by the SNMP manager. These MIBs provide information the SNMP manager needs to interpret the SNMP trap, event, and query messages sent by the FortiGate unit SNMP agent.

FortiGate core MIB files are available for download by going to **System > SNMP** and selecting the download link on the page.

The Fortinet implementation of SNMP includes support for most of RFC 2665 (Ethernet-like MIB) and most of RFC 1213 (MIB II). For more information, see “[Fortinet MIBs](#)”. RFC support for SNMP v3 includes Architecture for SNMP Frameworks (RFC 3411), and partial support of User-based Security Model (RFC 3414).

SNMP traps alert you to events that occur such as a full log disk or a virus detected.

SNMP fields contain information about the FortiGate unit, such as CPU usage percentage or the number of sessions. This information is useful for monitoring the condition of the unit on an ongoing basis and to provide more information when a trap occurs.

The FortiGate SNMP v3 implementation includes support for queries, traps, authentication, and privacy. Authentication and encryption are configured in the CLI. See the `system snmp user` command in the FortiGate CLI Reference.

New SNMP trap for bypass events

When bypass mode is enabled or disabled on FortiGate units that are equipped with bypass interfaces and support AMC modules, a new SNMP trap is generated and logs bypass events.

Implement SNMP support for NAT Session monitoring which includes new SNMP OIDs

FortiOS 5.6 implements a new feature providing SNMP support for NAT session monitoring. The resulting new SNMP object identifier (OID) is:

```
FORTINET-FORTIGATE-  
MIB:fortinet.fnFortiGateMib.fgFirewall.fgFwIppools.fgFwIppTables.fgFwIppStatsTable.fgFwIppStatsEntry  
1.3.6.1.4.1.12356.101.5.3.2.1.1
```

Additionally, there are eight new items:

```
.fgFwIppStatsName .1  
.fgFwIppStatsType .2  
.fgFwIppStatsStartIp .3  
.fgFwIppStatsEndIp .4  
.fgFwIppStatsTotalSessions .5  
.fgFwIppStatsTcpSessions .6  
.fgFwIppStatsUdpSessions .7  
.fgFwIppStatsOtherSessions .8
```

SNMP configuration settings

Before a remote SNMP manager can connect to the FortiGate agent, you must configure one or more FortiGate interfaces to accept SNMP connections by going to **Network > Interfaces**. Select the interface and, in the **Administrative Access**, select **SNMP**.

For VDOMS, SNMP traps can only be sent on interfaces in the management VDOM. Traps cannot be sent over other interfaces outside the management VDOM.

To configure SNMP settings, go to **System > SNMP**.

SNMP Agent	Select to enable SNMP communication.
Description	Enter descriptive information about the FortiGate unit. The description can be up to 35 characters.
Location	Enter the physical location of the FortiGate unit. The system location description can be up to 35 characters long.
Contact	Enter the contact information for the person responsible for this FortiGate unit. The contact information can be up to 35 characters.
SNMP v1/v2c section To create a new SNMP community, see SNMP Community .	
Community Name	The name to identify the community.
Queries	Indicates whether queries protocols (v1 and v2c) are enabled or disabled. A green check mark indicates queries are enabled; a gray x indicates queries are disabled. If one query is disabled and another one enabled, there will still be a green check mark.
Traps	Indicates whether trap protocols (v1 and v2c) are enabled or disabled. A green check mark indicates traps are enabled; a gray x indicates traps are disabled. If one query is disabled and another one enabled, there will still be a green check mark.
Enable	Select the check box to enable or disable the community.
SNMP v3 section To create a new SNMP community, see SNMP Community page.	
User Name	The name of the SNMPv3 user
Security Level	The security level of the user
Notification Host	The IP address or addresses of the host
Queries	Indicates whether queries are enabled or disabled. A green check mark indicates queries are enabled; a gray x indicates queries are disabled
New SNMP Community page	
Community Name	Enter a name to identify the SNMP community
Hosts (section)	

IP Address	<p>Enter the IP address and Identify the SNMP managers that can use the settings in this SNMP community to monitor the FortiGate unit.</p> <p>You can also set the IP address to 0.0.0.0 to so that any SNMP manager can use this SNMP community.</p>
Delete	Removes an SNMP manager from the list within the Hosts section
Add	Select to add a blank line to the Hosts list. You can add up to eight SNMP managers to a single community.
Queries (section)	
Protocol	The SNMP protocol. In the v1 row, this means that the settings are for SNMP v1. In the v2c row, this means that the settings are for SNMP v2c.
Port	<p>Enter the port number (161 by default) that the SNMP managers in this community use for SNMP v1 and SNMP v2c queries to receive configuration information from the FortiGate unit. Select the Enable check box to activate queries for each SNMP version.</p> <p>Note: The SNMP client software and the FortiGate unit must use the same port for queries.</p>
Enable	Select to enable that SNMP protocol.
Traps (section)	
Protocol	The SNMP protocol. In the v1 row, this means that the settings are for SNMP v1. In the v2c row, this means that the settings are for SNMP v2c.
Local	<p>Enter the remote port numbers (port 162 for each by default) that the FortiGate unit uses to send SNMP v1 or SNMP v2c traps to the SNMP managers in this community. Select the Enable check box to activate traps for each SNMP version.</p> <p>Note: The SNMP client software and the FortiGate unit must use the same port for traps.</p>
Remote	<p>Enter the remote port number (port 162 is default) that the FortiGate unit uses to send SNMP v1 or v2c traps to the SNMP managers in this community.</p> <p>Note: The SNMP client software and the FortiGate unit must use the same port for queries.</p>
Enable	Select to activate traps for each SNMP version.

SNMP Event	<p>Enable each SNMP event for which the FortiGate unit should send traps to the SNMP managers in this community.</p> <p>CPU Over usage traps sensitivity is slightly reduced, by spreading values out over 8 polling cycles. This prevents sharp spikes due to CPU intensive short-term events such as changing a policy.</p> <p>Power Supply Failure event trap is available only on some models.</p> <p>AMC interfaces enter bypass mode event trap is available only on models that support AMC modules.</p>
Enable	Select to enable the SNMP event.
Create New SNMP V3 User	
User Name	Enter the name of the user.
Security Level	Select the type of security level the user will have.
Notification Host	Enter the IP address of the notification host. If you want to add more than one host, after entering the IP address of the first host, select the plus sign to add another host.
Enable Query	Select to enable or disable the query. By default, the query is enabled.
Port	Enter the port number in the field.
Events	Select the SNMP events that will be associated with that user.

Gigabit interfaces

When determining the interface speed of a FortiGate unit with a 10G interface, the IF-MIB.ifSpeed may not return the correct value. IF-MIB.ifSpeed is a 32-bit gauge used to report interface speeds in bits/second and cannot convert to a 64-bit value. The 32-bit counter wrap the output too fast to be accurate.

In this case, you can use the value ifHighSpeed. It reports interface speeds in megabits/second. This ensures that 10Gb interfaces report the correct value.

SNMP agent

You need to first enter information and enable the FortiGate SNMP Agent. Enter information about the FortiGate unit to identify it so that when your SNMP manager receives traps from the FortiGate unit, you will know which unit sent the information.

To configure the SNMP agent - GUI

1. Go to **System > SNMP**.
2. Select **Enable** for the **SNMP Agent**.
3. Enter a descriptive name for the agent.
4. Enter the location of the FortiGate unit.

5. Enter a contact or administrator for the SNMP Agent or FortiGate unit.
6. Select **Apply**.

To configure SNMP agent - CLI

```
config system snmp sysinfo
  set status enable
  set contact-info <contact_information>
  set description <description_of_FortiGate>
  set location <FortiGate_location>
end
```

SNMP community

An SNMP community is a grouping of devices for network administration purposes. Within that SNMP community, devices can communicate by sending and receiving traps and other information. One device can belong to multiple communities, such as one administrator terminal monitoring both a firewall SNMP and a printer SNMP community.

Add SNMP communities to your FortiGate unit so that SNMP managers can connect to view system information and receive SNMP traps.

You can add up to three SNMP communities. Each community can have a different configuration for SNMP queries and traps. Each community can be configured to monitor the FortiGate unit for a different set of events. You can also add the IP addresses of up to 8 SNMP managers to each community.

When the FortiGate unit is in virtual domain mode, SNMP traps can only be sent on interfaces in the management virtual domain. Traps cannot be sent over other interfaces.

To add an SNMP v1/v2c community - GUI

1. Go to **System > SNMP**.
2. In the **SNMP v1/v2c** area, select **Create New**.
3. Enter a **Community Name**.
4. Enter the IP address and Identify the SNMP managers that can use the settings in this SNMP community to monitor the FortiGate unit.
5. Select the interface if the SNMP manager is not on the same subnet as the FortiGate unit.
6. Enter the **Port** number that the SNMP managers in this community use for SNMP v1 and SNMP v2c queries to receive configuration information from the FortiGate unit. Select the **Enable** check box to activate queries for each SNMP version.
7. Enter the Local and Remote port numbers that the FortiGate unit uses to send SNMP v1 and SNMP v2c traps to the SNMP managers in this community.
8. Select the **Enable** check box to activate traps for each SNMP version.
9. Select **OK**.

To add an SNMP v1/v2c community - CLI

```
config system snmp community
  edit <index_number>
    set events <events_list>
    set name <community_name>
    set query-v1-port <port_number>
    set query-v1-status {enable | disable}
```

```

    set query-v2c-port <port_number>
    set query-v2c-status {enable | disable}
    set status {enable | disable}
    set trap-v1-lport <port_number>
    set trap-v1-rport <port_number>
    set trap-v1-status {enable | disable}
    set trap-v2c-lport <port_number>
    set trap-v2c-rport <port_number>
    set trap-v2c-status {enable | disable}
end

```

To add an SNMP v3 community - GUI

1. Go to **System > SNMP**.
2. In the **SNMP v3** area, select **Create New**.
3. Enter a **User Name**.
4. Select a **Security Level** and associated authorization algorithms.
5. Enter the IP address of the **Notification Host** SNMP managers that can use the settings in this SNMP community to monitor the FortiGate unit.
6. Enter the **Port** number that the SNMP managers in this community use to receive configuration information from the FortiGate unit. Select the **Enable** check box to activate queries for each SNMP version.
7. Select the **Enable** check box to activate traps.
8. Select **OK**.

To add an SNMP v3 community - CLI

```

config system snmp user
  edit <index_number>
    set security-level [auth-priv | auth-no-priv | no-auth-no-priv]
    set queries enable
    set query-port <port_number>
    set notify-hosts <ip_address>
    set events <event_selections>
end

```

Enabling on the interface

Before a remote SNMP manager can connect to the FortiGate agent, you must configure one or more FortiGate interfaces to accept SNMP connections.

To configure SNMP access - GUI

1. Go to **Network > Interfaces**.
2. Choose an interface that an SNMP manager connects to and select **Edit**.
3. In **Administrative Access**, select **SNMP**.
4. Select **OK**.

To configure SNMP access - CLI

```

config system interface
  edit <interface_name>
    set allowaccess snmp
end

```



If the interface you are configuring already has protocols that are allowed access, use the command `append allowaccess snmp` instead, or else the other protocols will be replaced. For more information, see Adding and removing options from lists.

Fortinet MIBs

The FortiGate SNMP agent supports Fortinet proprietary MIBs as well as standard RFC 1213 and RFC 2665 MIBs. RFC support includes support for the parts of RFC 2665 (Ethernet-like MIB) and the parts of RFC 1213 (MIB II) that apply to FortiGate unit configuration.

There are two MIB files for FortiGate units - the Fortinet MIB, and the FortiGate MIB. The Fortinet MIB contains traps, fields and information that is common to all Fortinet products. The FortiGate MIB contains traps, fields and information that is specific to FortiGate units. Each Fortinet product has its own MIB. If you use other Fortinet products you will need to download their MIB files as well. Both MIB files are used for FortiOS and FortiOS Carrier; there are no additional traps for the Carrier version of the operating system.

The Fortinet MIB and FortiGate MIB along with the two RFC MIBs are listed in tables in this section. To download the two FortiGate MIB files, visit the [Fortinet Support](#) website. The Fortinet MIB contains information for Fortinet products in general. the Fortinet FortiGate MIB includes the system information for the FortiGate and version of FortiOS. Both files are required for proper SNMP data collection.

To download the MIB files, go to **System > SNMP** and select a MIB link in the **FortiGate SNMP MIB** section.

Your SNMP manager may already include standard and private MIBs in a compiled database that is ready to use. You must add the Fortinet proprietary MIB to this database to have access to the Fortinet specific information.



There were major changes to the MIB files between FortiOS Carrier v3.0 and v4.0. You need to use the new MIBs for FortiOS Carrier v4.0 or you may mistakenly access the wrong traps and fields.

MIB files are updated for each version of FortiOS. When upgrading the firmware ensure that you updated the Fortinet FortiGate MIB file as well.

Fortinet MIBs

MIB file name or RFC	Description
FORTINET-CORE-MIB.mib	<p>The Fortinet MIB includes all system configuration information and trap information that is common to all Fortinet products.</p> <p>Your SNMP manager requires this information to monitor FortiGate unit configuration settings and receive traps from the FortiGate SNMP agent.</p>
FORTINET-FORTIGATE-MIB.mib	<p>The FortiGate MIB includes all system configuration information and trap information that is specific to FortiGate units.</p> <p>Your SNMP manager requires this information to monitor FortiGate configuration settings and receive traps from the FortiGate SNMP agent. FortiManager systems require this MIB to monitor FortiGate units.</p>

MIB file name or RFC	Description
RFC-1213 (MIB II)	<p>The FortiGate SNMP agent supports MIB II groups with these exceptions.</p> <ul style="list-style-type: none"> • No support for the EGP group from MIB II (RFC 1213, section 3.11 and 6.10). • Protocol statistics returned for MIB II groups (IP/ICMP/TCP/UDP/etc.) do not accurately capture all FortiGate traffic activity. More accurate information can be obtained from the information reported by the Fortinet MIB.
RFC-2665 (Ethernet-like MIB)	<p>The FortiGate SNMP agent supports Ethernet-like MIB information. FortiGate SNMP does not support for the dot3Tests and dot3Errors groups.</p>



SNMP improvements for dynamic routing include support for RFC 4750 OSPF Version 2 Management Information Base and RFC 5643 Management Information Base for OSPFv3. These changes add the capability of logging dynamic routing activity. Examples include sending OSPF routing events or changes to a syslog server or FortiAnalyzer or changes in neighborhood status.

Device detection for SNMP traps

This setting is related to the device detection feature. It allows SNMP traps to detect when a new device comes online. Within SNMP configurations there is a configurable timeout setting that periodically checks for the device. When a check determines that the device is present a trap is sent.

In the GUI, when configuring an SNMP object, one of the settings is a checkbox, under **SNMP Events** for **Device detected**.

To configure the SNMP object in the CLI use the following syntax:

```
config system snmp community
  edit <community ID number>
    set name <string>
    set events device-new
  end
```

In order to configure the idle timeout for the device, use the following syntax in the CLI:

```
config system global
  set device-idle-timeout <integer of time in seconds>
end
```

The time value for the field can be set from 30 to 31536000.

SNMP get command syntax

Normally, to get configuration and status information for a FortiGate unit, an SNMP manager would use an SNMP get commands to get the information in a MIB field. The SNMP get command syntax would be similar to:

```
snmpget -v2c -c <community_name> <address_ipv4> {<OID> | <MIB_field>}
```

...where...

`<community_name>` is an SNMP community name added to the FortiGate configuration. You can add more than one community name to a FortiGate SNMP configuration. The most commonly used community name is `public`.

`<address_ipv4>` is the IP address of the FortiGate interface that the SNMP manager connects to.

`{<OID> | <MIB_field>}` is the object identifier (OID) for the MIB field or the MIB field name itself.

The `SNMP get` command gets firmware version running on the FortiGate unit. The community name is `public`. The IP address of the interface configured for SNMP management access is `10.10.10.1`. The firmware version MIB field is `fgSysVersion` and the OID for this MIB field is `1.3.6.1.4.1.12356.101.4.1.1`. The first command uses the MIB field name and the second uses the OID:

```
snmpget -v2c -c public 10.10.10.1 fgSysVersion.0
snmpget -v2c -c public 10.10.10.1 1.3.6.1.4.1.12356.101.4.1.1.0
```



The OIDs and object names used in these examples are dependent on the version of MIB and are subject to change.

Replacement messages

The replacement message list in **System > Replacement Messages** enables you to view and customize replacement messages. Highlight the replacement messages you wish to edit and customize the message content to your requirements. Hit **Save** when done. If you do not see the message you want to edit, select the Extended View option (in the upper right-hand corner of the screen).

If you make a mistake, you can select the **Restore Default** to return to the original message and code base.

For connections requiring authentication, the FortiGate uses HTTP to send an authentication disclaimer page for the user to accept before a security policy is in effect. Therefore, you must initiate HTTP traffic first in order to trigger the authentication disclaimer page. Once the disclaimer is accepted, you can send whatever traffic is allowed by the security policy.

Replacement message images

You can add images to replacement messages to:

- Disclaimer pages
- Login pages
- Declined disclaimer pages
- Login failed page
- Login challenge pages
- Keepalive pages

Image embedding is also available to the endpoint NAC download portal and recommendation portal replacement messages, as well as HTTP replacement messages.

Supported image formats are GIF, JPEG, TIFF and PNG. The maximum file size supported is 6000 bytes.

Adding images to replacement messages

To upload an image for use in a message

1. Go to **System > Replacement Messages**.
2. Select **Manage Images** at the top of the page.
3. Select **Create New**.
4. Enter a **Name** for the image.
5. Select the **Content Type**.
6. Select **Browse** to locate the file and select **OK**.

The image that you include in a replacement message, must have the following html:

```
<img src=%%IMAGE: <config_image_name>%% size=<bytes> >
```

For example:

```
<img src=%%IMAGE: logo_hq%% size=4272>
```

Modifying replacement messages

Replacement messages can be modified to include a message or content that suits your organization.

Use the expand arrows to view the replacement message list for a given category. Messages are in HTML format. To change a replacement message, go to **System > Replacement Messages** and select the replacement message that you want to modify. At the bottom pane of the window, you can see the message on one side and the HTML code on the other side. The message view changes in real-time as you change the content.

A list of common replacement messages appears in the main window. To see the entire list and all categories of replacement messages, in the upper-right corner of the window, select **Extended View**.

Replacement message categories

Alert email replacement messages

The FortiGate unit adds the replacement messages listed in this category to alert email messages sent to administrators. If you enable the option **Send alert email for logs based on severity** in **Log & Report**, you control whether or not replacement messages are sent by alert email based on how you set the **Minimum log level**.

For more information on Alert emails, see the Monitoring chapter.

Authentication replacement messages

The FortiGate unit uses the text of the authentication replacement messages for various user authentication HTML pages that are displayed when a user is required to authenticate because a security policy includes at least one identity-based policy that requires firewall users to authenticate.

These replacement message pages are for authentication using HTTP and HTTPS. You cannot customize the firewall authentication messages for FTP and Telnet.

The authentication login page and the authentication disclaimer include replacement tags and controls not found on other replacement messages.

Users see the authentication login page when they use a VPN or a security policy that requires authentication. You can customize this page in the same way as you modify other replacement messages.

There are some unique requirements for these replacement messages:

- The login page must be an HTML page containing a form with `ACTION="/"` and `METHOD="POST"`
- The form must contain the following hidden controls:
 - `<INPUT TYPE="hidden" NAME="%%MAGICID%%" VALUE="%%MAGICVAL%%">`
 - `<INPUT TYPE="hidden" NAME="%%STATEID%%" VALUE="%%STATEVAL%%">`
 - `<INPUT TYPE="hidden" NAME="%%REDIRID%%" VALUE="%%PROTURI%%">`
- The form must contain the following visible controls:
 - `<INPUT TYPE="text" NAME="%%USERNAMEID%%" size=25>`
 - `<INPUT TYPE="password" NAME="%%PASSWORDID%%" size=25>`

Example

The following is an example of a simple authentication page that meets the requirements listed above.

```
<HTML><HEAD><TITLE>Firewall Authentication</TITLE></HEAD>
  <BODY><H4>You must authenticate to use this service.</H4>
  <FORM ACTION="/" method="post">
    <INPUT NAME="%%MAGICID%%" VALUE="%%MAGICVAL%%" TYPE="hidden">
  <TABLE ALIGN="center" BGCOLOR="#00cccc" BORDER="0"
    CELLPADDING="15" CELLSPACING="0" WIDTH="320"><TBODY>
```

```
<TR><TH>Username:</TH>
  <TD><INPUT NAME="%%USERNAMEID%%" SIZE="25" TYPE="text"> </TD></TR>
<TR><TH>Password:</TH>
  <TD><INPUT NAME="%%PASSWORDID%%" SIZE="25" TYPE="password"> </TD></TR>
<TR><TD COLSPAN="2" ALIGN="center" BGCOLOR="#00cccc">
  <INPUT NAME="%%STATEID%%" VALUE="%%STATEVAL%%" TYPE="hidden">
    <INPUT NAME="%%REDIRID%%" VALUE="%%PROTURI%%" TYPE="hidden">
    <INPUT VALUE="Continue" TYPE="submit"> </TD></TR>
</TBODY></TABLE></FORM></BODY></HTML>
```

Captive Portal Default replacement messages

The Captive Portal Default replacement messages are used for wireless authentication only. You must have a VAP interface with the security set as captive portal to trigger these replacement messages.

Device Detection Portal replacement message

The FortiGate unit displays the replacement message when the FortiGate unit cannot determine the type of BYOD or handheld device is used to connect the network.

Email replacement messages

The FortiGate unit sends the mail replacement messages to email clients using IMAP, POP3, or SMTP when an event occurs such as antivirus blocking a file attached to an email that contains a virus. Email replacement messages are text messages.

If the FortiGate unit supports SSL content scanning and inspection these replacement messages can also be added to IMAPS, POP3S, and SMTPS email messages.

Endpoint Control replacement message

The FortiGate unit displays the replacement message when the FortiClient Endpoint Security software is not installed or registered correctly with the FortiGate unit.

FortiGuard Web Filtering replacement messages

The FortiGate unit sends the FortiGuard Web Filtering replacement messages to web browsers using the HTTP protocol when FortiGuard web filtering blocks a URL, provides details about blocked HTTP 4xx and 5xx errors, and for FortiGuard overrides. FortiGuard Web Filtering replacement messages are HTTP pages.

If the FortiGate unit supports SSL content scanning and inspection and if **Protocol Recognition > HTTPS Content Filtering Mode** is set to Deep Scan in the antivirus profile, these replacement messages can also replace web pages downloaded using the HTTPS protocol.

FTP replacement messages

The FortiGate unit sends the FTP replacement messages listed in the table below to FTP clients when an event occurs such as antivirus blocking a file that contains a virus in an FTP session. FTP replacement messages are text messages.

HTTP replacement messages

The FortiGate unit sends the HTTP replacement messages listed in the following table to web browsers using the HTTP protocol when an event occurs such as antivirus blocking a file that contains a virus in an HTTP session.

HTTP replacement messages are HTML pages.

If the FortiGate unit supports SSL content scanning and inspection, and if under HTTPS in the protocol option list has Enable Deep Scan enabled, these replacement messages can also replace web pages downloaded using the HTTPS protocol.

NNTP replacement messages

The FortiGate unit sends the NNTP replacement messages NNTP clients when an event occurs such as antivirus blocking a file attached to an NNTP message that contains a virus. NNTP replacement messages are text messages.

Spam replacement messages

The FortiGate unit adds the Spam replacement messages to SMTP server responses if the email message is identified as spam and the spam action is discard. If the FortiGate unit supports SSL content scanning and inspection these replacement messages can also be added to SMTPS server responses.

NAC quarantine replacement messages

The page that is displayed for the user depends on whether NAC quarantine blocked the user because a virus was found, a DoS sensor detected an attack, an IPS sensor detected an attack, or a DLP rule with action set to **Quarantine IP address** or **Quarantine Interface** matched a session from the user.

The default messages inform the user of why they are seeing this page and recommend they contact the system administrator. You can customize the pages as required, for example to include an email address or other contact information or if applicable a note about how long the user can expect to be blocked.

SSL VPN replacement messages

The SSL VPN login replacement message is an HTML replacement message that formats the FortiGate SSL VPN portal login page. You can customize this replacement message according to your organization's needs. The page is linked to FortiGate functionality and you must construct it according to the following guidelines to ensure that it will work.

- The login page must be an HTML page containing a form with `ACTION="%%SSL_ACT%%"` and `METHOD="%%SSL_METHOD%%"`
- The form must contain the `%%SSL_LOGIN%%` tag to provide the login form.
- The form must contain the `%%SSL_HIDDEN%%` tag.

Web Proxy replacement messages

The FortiGate unit sends Web Proxy replacement messages when a web proxy event occurs that is detected and matches the web proxy configuration. These replacement messages are web pages that appear within your web browser.

The following web proxy replacement messages require an identity-based security policy so that the web proxy is successful. You can also enable FTP-over-HTTP by selecting the **FTP** option in **System > Network > Explicit Proxy**.

Traffic quota control replacement messages

When user traffic is going through the FortiGate unit and it is blocked by traffic shaping quota controls, users see the **Traffic shaper block message** or the **Per IP traffic shaper block message** when they attempt to connect through the FortiGate unit using HTTP.

The traffic quota HTTP pages should contain the `%%QUOTA_INFO%%` tag to display information about the traffic shaping quota setting that is blocking the user.

MM1 replacement messages

MM1 replacement messages are sent when, during MMS content scanning, FortiOS Carrier detects, for example a virus, using the MMS profile.

You must have **Remove Blocked** selected within the MMS profile if you want to remove the content that is intercepted during MMS scanning on the FortiGate unit.

MM3 replacement messages

MM3 replacement messages are sent when, during MMS content scanning, FortiOS Carrier detects, for example a virus, using the MMS profile.

You must have **Remove Blocked** selected within the MMS profile if you want to remove the content that is intercepted during MMS scanning on the unit.

MM4 replacement messages

MM4 replacement messages are sent when, during MMS content scanning, FortiOS Carrier detects, for example a virus, using the MMS profile.

MM7 replacement messages

MM7 replacement messages are sent when, during MMS content scanning, FortiOS Carrier detects, for example a virus, using the MMS profile.

MMS replacement messages

The MMS replacement message is sent when a section of an MMS message has been replaced because it contains a blocked file. This replacement message is in HTML format.

The message text is:

```
<HTML><BODY>This section of the message has been replaced because it contained a blocked
file</BODY></HTML>
```

Replacement message groups

Replacement message groups enable you to view common messages in groups for large carriers. Message groups can be configured by going to **Config > Replacement Message Group**.

Using the defined groups, you can manage specific replacement messages from a single location, rather than searching through the entire replacement message list.

If you enable virtual domains (VDOMs) on the FortiGate unit, replacement message groups are configured separately for each virtual domain. Each virtual domain has its own default replacement message group, configured from **System > Replacement Messages Group**.

When you modify a message in a replacement message group, a Reset icon appears beside the message in the group. You can select this Reset icon to reset the message in the replacement message group to the default version.

All MM1/4/7 notification messages for FortiOS Carrier (and MM1 retrieve-conf messages) can contain a SMIL layer and all MM4 notification messages can contain an HTML layer in the message. These layers can be used to brand messages by using logos uploaded to the FortiGate unit via the 'Manage Images' link found on the replacement message group configuration page.

Administration for schools

For a system administrator in a school system, it is difficult to maintain a network and access to the Internet. There are potential legal liabilities if content is not properly filtered and students gain access to pornography and other non-productive and potentially dangerous content. This section describes some basic practices administrators can employ.

Security policies

The default security policies in FortiOS allow all traffic on all ports and all IP addresses. While applying security profiles can help to block viruses, detect attacks and prevent spam, this doesn't provide a solid overall security option. The best approach is a layered approach; the first layer being the security policy.

When creating outbound security policies, you need to identify what the students are allowed to do. Generally, they surf the web, connect to FTP sites, send and receive email, and so on.

Once you know what the students need to do, you can research the software used and determine the ports the applications use. For example, if the students only require web surfing, then there are only two ports (80 - HTTP and 443 - HTTPS) needed to complete their tasks. Setting the security policies to only allow traffic through two ports (rather than all 65,000), this will significantly lower any possible exploits. By restricting the ports to known services, means stopping the use of proxy servers, as many of them operate on a non-standard port to hide their traffic from URL filtering or HTTP inspection.

DNS

You should restrict the students use of DNS. We recommend that you point to an internal DNS server and only allow those devices out on port 53.

If there is no internal DNS server, then you should establish a restrictive list of allowed DNS servers students can use. One possible exploit would be for them to set up their own DNS server at home that serves different IPs for known hosts, such as having Google.com sent back the IP for playboy.com.

Encrypted traffic (HTTPS)

Generally speaking, students should not be allowed to access encrypted web sites. Encrypted traffic can't be sniffed, and therefore, can't be monitored. HTTPS traffic should only be allowed when necessary. Most web sites a student needs to access are HTTP, not HTTPS. Due to the nature of HTTPS protocol, and the fact that encryption is an inherent security risk to your network, its use should be restricted.

To ensure that students only visit HTTPS sites required for schoolwork, we recommend that you add a security policy that designates a list of allowed secure sites.

FTP

For the most part, students should not be using FTP. FTP is not HTTP or HTTPS so you can't use URL filtering to restrict where they go. This can be controlled with destination IPs in the security policy. With a policy that specifically outlines which FTP addresses are allowed, all other will be blocked.

Example security policies

Given these recommended practices, a set of security policies could look like the following illustration. In a large setup, all the IPs for the students are treated by one of these four policies.

Simple security policy setup

ID	Name	From	To	Source	Destination	Schedule	Service	Action	Status
2	restrict https	internal	wan1	all Students	Allowed Websites	always	HTTPS	✓ ACCEPT	✓ Enabled
3	all http	internal	wan1	all Students	all	always	HTTP	✓ ACCEPT	✓ Enabled
4	allowed dns	internal	wan1	all Students	Allowed DNS	always	DNS	✓ ACCEPT	✓ Enabled
5	allowed ftp	internal	wan1	all Students	Allowed FTP	always	FTP	✓ ACCEPT	✓ Enabled
0	Implicit Deny	any	any	all	all	always	ALL	✗ DENY	

The last policy in the list is the deny policy that is configured by default. The deny policy ensures that any traffic making it to this point is stopped. It can also help in further troubleshooting by viewing the logs for denied traffic.

With these policies in place, even before packet inspection occurs, the FortiGate, and the network are fairly secure. Should any of the UTM profiles fail, there is still a basic level of security.

Security profiles

AntiVirus profiles

You should enable antivirus screening for any service that is enabled in the security policies. In the case above, HTTP, FTP, as well as POP3 and SMTP (assuming there is email access for students). There is not a virus scan option for HTTPS, because the content is encrypted. Generally speaking, most of the network traffic will be students surfing the web.

To configure antivirus profiles in the GUI, go to Security Profiles > AntiVirus.

Web filtering

You should configure your FortiGate to filter of URLs, sites and content, should be performed by FortiGuard. It is easier for the network administrator. Web sites are constantly being monitored, and new ones reviewed and added to the FortiGuard databases every day. The FortiGuard categories provide an extensive list of offensive and non-productive sites.

There are additional settings to include in a web filtering profile to best contain a student's web browsing.

- Web URL filtering should be enabled to set up exemptions for web sites that are blocked or reasons other than category filtering. It also prevents the use of IP addresses to get around web filtering.
- Block invalid URLs - HTTPS only. This option inspects the HTTPS certificate and looks at the URL to ensure it's valid. It is common for proxy sites to create an HTTPS certificate with a garbage URL. If the site is legitimate, it should be set up correctly. If the site approach to security is to ignore it, then their security policy puts your network at risk and the site should be blocked.

- Enable **Block malicious URLs discovered by FortiSandbox**. If the FortiSandbox discovers a threat, the source URL will be added to the list of URLs to be blocked by the FortiGate.

To configure web filtering options, go to **Security Profiles > Web Filter**.

Categories and classifications

For the selection of what FortiGuard categories and classifications that should be blocked, that is purely based on the school system and its Internet information policy.

Email filtering

Other than specific teacher-led email inboxes, there is no reason why a student should be able to access, read or send personal email. Ports for POP3, SMTP and IMAP should not be opened in a security policies.

IPS

The intrusion prevention profiles are used to ensure the student PCs are not vulnerable to attacks and are not in a position to make attacks. As well, IPS can do more than simple vulnerability scans. With a FortiGuard subscription, IPS signatures are pushed to the FortiGate unit. New signatures are released constantly for various intrusions as they are discovered.

FortiOS includes a number of predefined IPS sensors that you can enable by default. To configure IPS sensors, go to **Security Profiles > Intrusion Prevention**.

Application control

Application control uses IPS signatures to limit the use of instant messaging and peer-to-peer applications which can lead to possible infections on a student's PC. FortiOS includes a number of pre-defined application categories. To configure and maintain application control profiles, go to **Security Profiles > Application Control**.

Some applications to consider include proxies, botnets, toolbars and P2P applications.

Logging

Turn on all logging. Every option in this section should be enabled. This is not where you decide what you are going to log. You are identifying what the profiles can log.

Logging everything is a way to monitor traffic on the network, see what student's are utilizing the most, and locate any potential holes in your security plan. Keeping this information may help to prove negligence later if necessary.

PPTP and L2TP

A virtual private network (VPN) is a way to use a public network, such as the Internet, as a vehicle to provide remote offices or individual users with secure access to private networks. FortiOS supports the Point-to-Point Tunneling Protocol (PPTP), which enables interoperability between FortiGate units and Windows or Linux PPTP clients. Because FortiGate units support industry standard PPTP VPN technologies, you can configure a PPTP VPN between a FortiGate unit and most third-party PPTP VPN peers.

This section describes how to configure PPTP and L2TP VPNs as well as PPTP passthrough.

This section includes the topics:

- [How PPTP VPNs work](#)
- [FortiGate unit as a PPTP server](#)
- [Configuring the FortiGate unit for PPTP VPN](#)
- [Configuring the FortiGate unit for PPTP pass through](#)
- [Testing PPTP VPN connections](#)
- [Logging VPN events](#)
- [Configuring L2TP VPNs](#)
- [L2TP configuration overview](#)

How PPTP VPNs work

The Point-to-Point Tunneling Protocol enables you to create a VPN between a remote client and your internal network. Because it is a Microsoft Windows standard, PPTP does not require third-party software on the client computer. As long as the ISP supports PPTP on its servers, you can create a secure connection by making relatively simple configuration changes to the client computer and the FortiGate unit.

PPTP uses Point-to-Point protocol (PPP) authentication protocols so that standard PPP software can operate on tunneled PPP links. PPTP packages data in PPP packets and then encapsulates the PPP packets within IP packets for transmission through a VPN tunnel.

When the FortiGate unit acts as a PPTP server, a PPTP session and tunnel is created as soon as the PPTP client connects to the FortiGate unit. More than one PPTP session can be supported on the same tunnel. FortiGate units support PAP, CHAP, and plain text authentication. PPTP clients are authenticated as members of a user group.


Traffic from one PPTP peer is encrypted using PPP before it is encapsulated using Generic Routing Encapsulation (GRE) and routed to the other PPTP peer through an ISP network. PPP packets from the remote client are addressed to a computer on the private network behind the FortiGate unit. PPTP packets from the remote client are addressed to the public interface of the FortiGate unit. See the figure below.



PPTP control channel messages are not authenticated, and their integrity is not protected. Furthermore, encapsulated PPP packets are not cryptographically protected and may be read or modified unless appropriate encryption software such as Secure Shell (SSH) or Secure File Transfer Protocol (SFTP) is used to transfer data after the tunnel is established.

As an alternative, you can use encryption software such as Microsoft Point-to-Point Encryption (MPPE) to secure the channel. MPPE is built into Microsoft Windows clients and can be installed on Linux clients. FortiGate units support MPPE.

Packet encapsulation



Shown above, traffic from the remote client is addressed to a computer on the network behind the FortiGate unit. When the PPTP tunnel is established, packets from the remote client are encapsulated and addressed to the FortiGate unit. The FortiGate unit forwards disassembled packets to the computer on the internal network.

When the remote PPTP client connects, the FortiGate unit assigns an IP address from a reserved range of IP addresses to the client PPTP interface. The PPTP client uses the assigned IP address as its source address for the duration of the connection.

When the FortiGate unit receives a PPTP packet, the unit disassembles the PPTP packet and forwards the packet to the correct computer on the internal network. The security policy and protection profiles on the FortiGate unit ensure that inbound traffic is screened and processed securely.

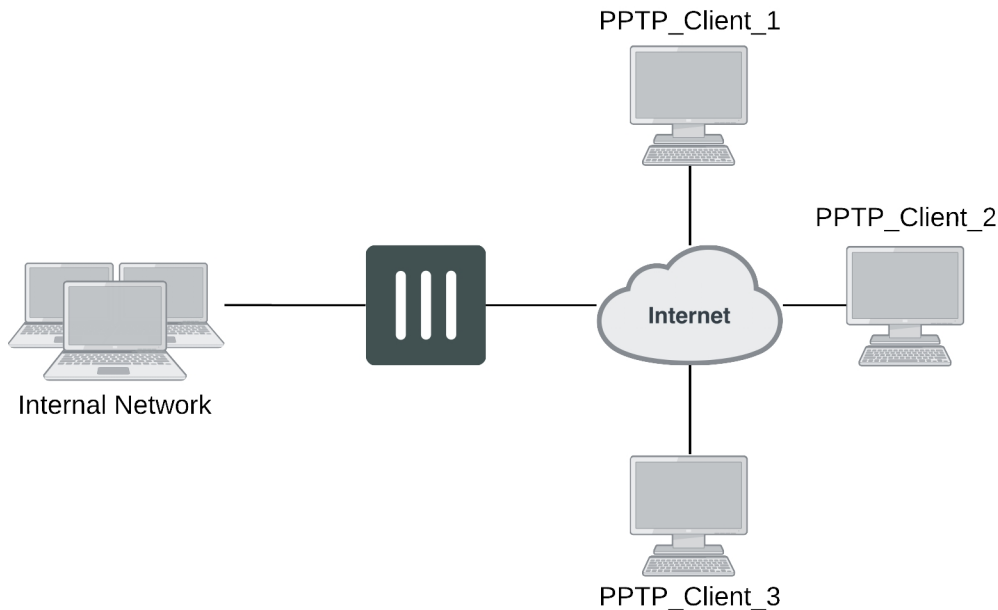


PPTP clients must be authenticated before a tunnel is established. The authentication process relies on FortiGate user group definitions, which can optionally use established authentication mechanisms such as RADIUS or LDAP to authenticate PPTP clients. All PPTP clients are challenged when a connection attempt is made.

FortiGate unit as a PPTP server

In the most common Internet scenario, the PPTP client connects to an ISP that offers PPP connections with dynamically-assigned IP addresses. The ISP forwards PPTP packets to the Internet, where they are routed to the FortiGate unit.

FortiGate unit as a PPTP server



If the FortiGate unit will act as a PPTP server, there are a number of steps to complete:

- Configure user authentication for PPTP clients.
- Enable PPTP.
- Specify the range of addresses that are assigned to PPTP clients when connecting
- Configure the security policy.

Configuring user authentication for PPTP clients

To enable authentication for PPTP clients, you must create user accounts and a user group to identify the PPTP clients that need access to the network behind the FortiGate unit. Within the user group, you must add a user for each PPTP client.

You can choose to use a plain text password for authentication or forward authentication requests to an external RADIUS, LDAP, or TACACS+ server. If password protection will be provided through a RADIUS, LDAP, or TACACS+ server, you must configure the FortiGate unit to forward authentication requests to the authentication server.

This example creates a basic user/password combination.

Configuring a user account

To add a local user - GUI

1. Go to **User & Device > User Definition** and select **Create New**.
2. Select **Local User**
3. Enter a **User Name**.
4. Enter a **Password** for the user. The password should be at least six characters.
5. Select **OK**.

To add a local user - CLI

```
config user local
  edit <username>
    set type password
    set passwd <password>
  end
```

Configuring a user group

To ease configuration, create user groups that contain users in similar categories or departments.

To create a user group - GUI

1. Go to **User & Device > User Group** and select **Create New**.
2. Enter a **Name** for the group.
3. Select the **Type** of **Firewall**.
4. From the **Available Users** list, select the required users and select the right-facing arrow to add them to the **Members** list.
5. Select **OK**.

To create a user group - CLI

```
config user group
  edit <group_name>
    set group-type firewall
    set member <user_names>
  end
```

Enabling PPTP and specifying the PPTP IP address range

The PPTP address range specifies the range of addresses reserved for remote PPTP clients. When a PPTP client connects to the FortiGate unit, the client is assigned an IP address from this range. Afterward, the FortiGate unit uses the assigned address to communicate with the PPTP client.

The address range that you reserve can be associated with private or routable IP addresses. If you specify a private address range that matches a network behind the FortiGate unit, the assigned address will make the PPTP client appear to be part of the internal network.

PPTP requires two IP addresses, one for each end of the tunnel. The PPTP address range is the range of addresses reserved for remote PPTP clients. When the remote PPTP client establishes a connection, the

FortiGate unit assigns an IP address from the reserved range of IP addresses to the client PPTP interface or retrieves the assigned IP address from the PPTP user group. If you use the PPTP user group, you must also define the FortiGate end of the tunnel by entering the IP address of the unit in **Local IP** (web-based manager) **or** `local-ip` (CLI). The PPTP client uses the assigned IP address as its source address for the duration of the connection.

PPTP configuration is only available through the CLI. In the example below, PPTP is enabled with the use of an IP range of 192.168.1.1 to 192.168.1.10 for addressing and the user group is `hr_staff`.



FortiOS 5.4.0 and later versions allow the start and end IPs in the PPTP address range to be in the same 16-bit subnet. Earlier versions require that the start and end IPs in the PPTP address range be in the same 24-bit subnet, for example, 192.168.1.1 - 192.168.1.254. .

```
config vpn pptp
  set status enable
  set ip-mode range
  set eip 192.168.1.10
  set sip 192.168.1.1
  set usrgrp hr_staff
end
```

In this example, PPTP is enabled with the use of a user group for addressing, where the IP address of the PPTP server is 192.168.1.2 and the user group is `hr_admin`.

```
config vpn pptp
  set status enable
  set ip-mode range
  set local-ip 192.168.2.1
  set usrgrp hr_admin
end
```

Adding the security policy

The security policy specifies the source and destination addresses that can generate traffic inside the PPTP tunnel and defines the scope of services permitted through the tunnel. If a selection of services are required, define a service group.

To configure the firewall for the PPTP tunnel - GUI

1. Go to **Policy & Objects > IPv4** or **Policy & Objects > IPv6** and select **Create New**.
2. Complete the following and select **OK**:

Incoming Interface	The FortiGate interface connected to the Internet.
Source Address	Select the name that corresponds to the range of addresses that you reserved for PPTP clients.
Outgoing Interface	The FortiGate interface connected to the internal network.
Destination Address	Select the name that corresponds to the IP addresses behind the FortiGate unit.

Schedule	always
Service	ALL
Action	ACCEPT

To configure the firewall for the PPTP tunnel - CLI

```
config firewall policy or config firewall policy6
edit 1
    set srcintf <interface to internet>
    set dstintf <interface to internal network>
    set srcaddr <reserved_range>
    set dstaddr <internal_addresses>
    set action accept
    set schedule always
    set service ALL
end
```

Configuring the FortiGate unit for PPTP VPN

To arrange for PPTP packets to pass through the FortiGate unit to an external PPTP server, perform the following tasks in the order given:

- Configure user authentication for PPTP clients.
- Enable PPTP on the FortiGate unit and specify the range of addresses that can be assigned to PPTP clients when they connect.
- Configure PPTP pass through on the FortiGate unit.

Configuring the FortiGate unit for PPTP passthrough

To forward PPTP packets to a PPTP server on the network behind the FortiGate unit, you need to perform the following configuration tasks on the FortiGate unit:

- Define a virtual IP address that points to the PPTP server.
- Create a security policy that allows incoming PPTP packets to pass through to the PPTP server.



The address range is the external (public) ip address range which requires access to the internal PPTP server through the FortiGate virtual port-forwarding firewall.

IP addresses used in this document are fictional and follow the technical documentation guidelines specific to Fortinet. Real external IP addresses are not used.

Configuring a virtual IP address

The virtual IP address will be the address of the PPTP server host.

To define a virtual IP for PPTP passthrough - GUI

1. Go to **Policy & Objects > Virtual IPs**.
2. Select **Create New**.

3. Choose the **VIP Type**.
4. Enter the name of the VIP, for example, `PPTP_Server`.
5. Select the **External Interface** where the packets will be received for the PPTP server.
6. Enter the **External IP Address** for the VIP.
7. Select **Port Forwarding**.
8. Set the **Protocol to TCP**.
9. Enter the **External Service Port** of 1723, the default for PPTP.
10. Enter the **Map to Port** to 1723.
11. Select **OK**.

To define a virtual IP for PPTP passthrough - CLI

```
config firewall vip or config firewall vip6
edit PPTP_Server
    set extintf <interface>
    set extip <ip_address>
    set portforward enable
    set protocol tcp
    set extport 1723
    set mappedport 1723
    set mappedip <destination IP address range>
end
```

You can also use `config firewall vip46` to define a virtual IP from an IPv4 address to an IPv6 address or `config firewall vip64` to define a virtual IP from an IPv6 address to an IPv4 address.

Configuring a port-forwarding security policy

To create a port-forwarding security policy for PPTP passthrough you must first create an address range reserved for the PPTP clients.

To create an address range - GUI

1. Go to **Policy & Objects > Addresses** and select **Create New**.
2. Select a **Category**.
3. Enter a **Name** for the range, for example, `External_PPTP`.
4. Select a **Type** of **Subnet/IP Range**.
5. Enter the IP address range.
6. Select the **Interface** to the Internet.
7. Select **OK**.

To create an address range - CLI

```
config firewall address OR config firewall address6
edit External_PPTP
    set type ip_range
    set start-ip <ip_address>
    set end-ip <ip_address>
    set associated-interface <internet_interface>
end
```

With the address set, you can add the security policy.

To add the security policy - GUI

1. Go to **Policy & Objects > IPv4** or **Policy & Objects > IPv6** and select **Create New**.
2. Complete the following and select **OK**:

Incoming Interface	The FortiGate interface connected to the Internet.
Source Address	Select the address range created in the previous step.
Outgoing Interface	The FortiGate interface connected to the PPTP server.
Destination Address	Select the VIP address created in the previous steps.
Schedule	always
Service	PPTP
Action	ACCEPT

To add the security policy - CLI

```
config firewall policy or config firewall policy6
  edit <policy_number>
    set srcintf <interface to internet>
    set dstintf <interface to PPTP server>
    set srcaddr <address_range>
    set dstaddr <PPTP_server_address>
    set action accept
    set schedule always
    set service PPTP
  end
```

Testing PPTP VPN connections

To confirm that a PPTP VPN between a local network and a dialup client has been configured correctly, at the dialup client, issue a ping command to test the connection to the local network. The PPTP VPN tunnel initializes when the dialup client attempts to connect.

Logging VPN events

PPTP VPN, activity is logged when enabling VPN logging. The FortiGate unit connection events and tunnel status (up/down) are logged.

To log VPN events

1. Go to **Log & Report > Log Settings**.
2. Enable **Event Logging**.
3. Select **VPN activity event**.
4. Select **Apply**.

To view event logs

1. Go to **Log & Report > VPN Events**.
2. If the option is available to set the log location list, select disk or memory.
3. Select a log event and select **Details**.

Configuring L2TP VPNs

This section describes how to configure a FortiGate unit to establish a Layer Two Tunneling Protocol (L2TP) tunnel with a remote dialup client. The FortiGate implementation of L2TP enables a remote dialup client to establish an L2TP tunnel with the FortiGate unit directly.

According to RFC 2661, an Access Concentrator (LAC) can establish an L2TP tunnel with an L2TP Network Server (LNS). In a typical scenario, the LAC is managed by an ISP and located on the ISP premises; the LNS is the gateway to a private network. When a remote dialup client connects to the Internet through the ISP, the ISP uses a local database to establish the identity of the caller and determine whether the caller needs access to an LNS through an L2TP tunnel. If the services registered to the caller indicate that an L2TP connection to the LNS is required, the ISP LAC attempts to establish an L2TP tunnel with the LNS.

A FortiGate unit can be configured to act as an LNS. The FortiGate implementation of L2TP enables a remote dialup client to establish an L2TP tunnel with the FortiGate unit directly, bypassing any LAC managed by an ISP. The ISP must configure its network access server to forward L2TP traffic from the remote client to the FortiGate unit directly whenever the remote client requires an L2TP connection to the FortiGate unit.

When the FortiGate unit acts as an LNS, an L2TP session and tunnel is created as soon as the remote client connects to the FortiGate unit. The FortiGate unit assigns an IP address to the client from a reserved range of IP addresses. The remote client uses the assigned IP address as its source address for the duration of the connection.

More than one L2TP session can be supported on the same tunnel. FortiGate units can be configured to authenticate remote clients using a plain text user name and password, or authentication can be forwarded to an external RADIUS or LDAP server. L2TP clients are authenticated as members of a user group.

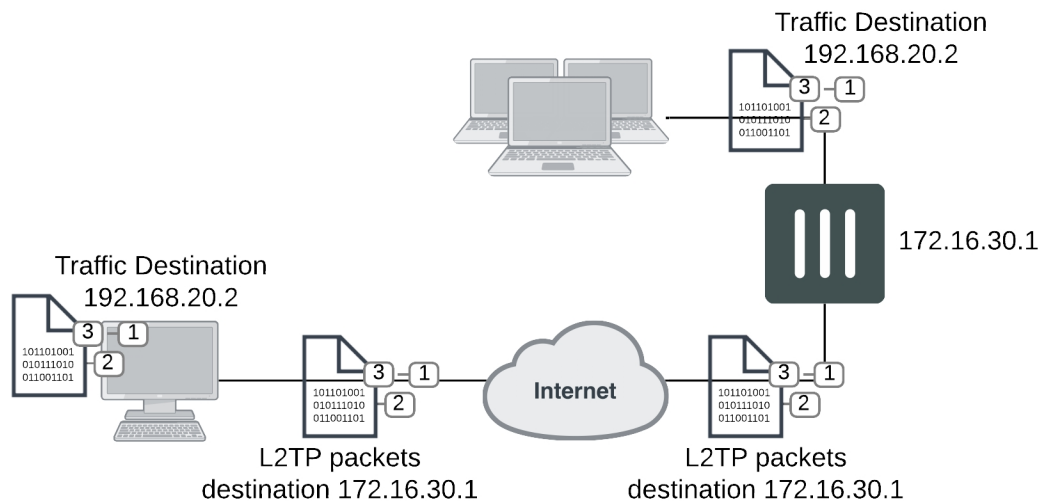


For site-to-site connections, Windows servers use IPsec encryption when you configure the VPN to connect to an L2TP server.

Traffic from the remote client must be encrypted using IPsec before it is encapsulated and routed to the FortiGate unit. Packets originating at the remote client are addressed to a computer on the private network behind the FortiGate unit. Encapsulated packets are addressed to the public interface of the FortiGate unit. See the figure below.

When the FortiGate unit receives an L2TP packet, the unit disassembles the packet and forwards the packet to the correct computer on the internal network. The security policy and protection profiles on the FortiGate unit ensure that inbound traffic is screened and processed securely.

L2TP encapsulation




FortiGate units cannot deliver non-IP traffic such as Frame Relay or ATM frames encapsulated in L2TP packets — FortiGate units support the IPv4 and IPv6 addressing schemes only

Network topology

The remote client connects to an ISP that determines whether the client requires an L2TP connection to the FortiGate unit. If an L2TP connection is required, the connection request is forwarded to the FortiGate unit directly.

Example L2TP configuration



L2TP infrastructure requirements

- The FortiGate unit must be operating in NAT mode and have a static public IP address.
- The ISP must configure its network access server to forward L2TP traffic from remote clients to the FortiGate unit directly.
- The remote client must not generate non-IP traffic (Frame Relay or ATM frames).

L2TP configuration overview

To configure a FortiGate unit to act as an LNS, you perform the following tasks:

- Create an L2TP user group containing one user for each remote client.
- Enable L2TP on the FortiGate unit and specify the range of addresses that can be assigned to remote clients when they connect.
- Define firewall source and destination addresses to indicate where packets transported through the L2TP tunnel will originate and be delivered.
- Create the security policy and define the scope of permitted services between the source and destination addresses.
- Configure the remote clients.

Authenticating L2TP clients

L2TP clients must be authenticated before a tunnel is established. The authentication process relies on FortiGate user group definitions, which can optionally use established authentication mechanisms such as RADIUS or LDAP to authenticate L2TP clients. All L2TP clients are challenged when a connection attempt is made.

To enable authentication, you must create user accounts and a user group to identify the L2TP clients that need access to the network behind the FortiGate unit.

You can choose to use a plain text password for authentication or forward authentication requests to an external RADIUS or LDAP server. If password protection will be provided through a RADIUS or LDAP server, you must configure the FortiGate unit to forward authentication requests to the authentication server.

Enabling L2TP and specifying an address range

The L2TP address range specifies the range of addresses reserved for remote clients. When a remote client connects to the FortiGate unit, the client is assigned an IP address from this range. Afterward, the FortiGate unit uses the assigned address to communicate with the remote client.

The address range that you reserve can be associated with private or routable IP addresses. If you specify a private address range that matches a network behind the FortiGate unit, the assigned address will make the remote client appear to be part of the internal network.

To enable L2TP and specify the L2TP address range, use the `config vpn l2tp` CLI command.

The following example shows how to enable L2TP and set the L2TP address range using a starting address of 192.168.10.80 and an ending address of 192.168.10.100 for an existing group of L2TP users named L2TP_users:

```
config vpn l2tp
  set sip 192.168.10.80
  set eip 192.168.10.100
  set status enable
  set usrgroup L2TP_users
end
```

Defining firewall source and destination addresses

Before you define the security policy, you must define the source and destination addresses of packets that are to be transported through the L2TP tunnel:

- For the source address, enter the range of addresses that you reserved for remote L2TP clients (for example 192.168.10.[80-100]).
- For the destination address, enter the IP addresses of the computers that the L2TP clients need to access on the private network behind the FortiGate unit (for example, 172.16.5.0/24 for a subnet, or 172.16.5.1 for a server or host, or 192.168.10.[10-15] for an IP address range).

To define the firewall source address

1. Go to **Policy & Objects > Addresses** and select **Create New**.
2. Select **Address**.
3. In the **Name** field, type a name that represents the range of addresses that you reserved for remote clients (for example, Ext_L2TPrange).
4. In **Type**, select **IP Range**.
5. In the **Subnet / IP Range** field, type the corresponding IP address range.
6. In **Interface**, select the FortiGate interface that connects to the clients.
7. This is usually the interface that connects to the Internet.
8. Select **OK**.

To define the firewall destination address

1. Go to **Policy & Objects > Addresses** and select **Create New**.
2. In the **Address Name** field, type a name that represents a range of IP addresses on the network behind the FortiGate unit (for example, `Int_L2TPaccess`).
3. In **Type**, select **IP Range**.
4. In the **IP Range** field, type the corresponding IP address range.
5. In **Interface**, select the FortiGate interface that connects to the network behind the FortiGate unit.
6. Select **OK**.

Adding the security policy

The security policy specifies the source and destination addresses that can generate traffic inside the L2TP tunnel and defines the scope of services permitted through the tunnel. If a selection of services are required, define a service group.

To define the traffic and services permitted inside the L2TP tunnel

1. Go to **Policy & Objects > IPv4** or **Policy & Objects > IPv6** and select **Create New**.
2. Enter these settings:

Name	Input a name for the policy.
Incoming Interface	Select the FortiGate interface to the Internet.
Outgoing Interface	Select the FortiGate interface to the internal (private) network.
Source Address	Select the name that corresponds to the address range that reserved for L2TP clients (for example, <code>Ext_L2TPrange</code>).
Destination Address	Select the name that corresponds to the IP addresses behind the FortiGate unit (for example, <code>Int_L2TPaccess</code>).
Schedule	Select ALWAYS, or if a select schedule is required instead, select a schedule that you defined previously.
Service	Select ALL, or if selected services are required instead, select the service group that you defined previously.
Action	ACCEPT

3. Select **OK**.

Configuring a Linux client

This procedure outlines how to install L2TP client software and run an L2TP tunnel on a Linux computer. Obtain an L2TP client package that meets your requirements (for example, `rp-l2tp`). If needed to encrypt traffic, obtain L2TP client software that supports encryption using IPsec.

To establish an L2TP tunnel with a FortiGate unit that has been set up to accept L2TP connections, you can obtain and install the client software following these guidelines:

1. If encryption is required, you will need to verify the IPsec configuration.
2. Download and install the L2TP client package.

3. Configure an L2TP connection to run the L2TP program.
4. Configure routes to determine whether all or some of your network traffic will be sent through the tunnel. You must define a route to the remote network over the L2TP link and a host route to the FortiGate unit.
5. Run `l2tpd` to start the tunnel.

Follow the software supplier's documentation to complete the steps.

To configure the system, you need to know the public IP address of the FortiGate unit, and the user name and password that has been set up on the FortiGate unit to authenticate L2TP clients. Contact the FortiGate administrator if required to obtain this information.

Monitoring L2TP sessions

You can display a list of all active sessions and view activity by port number. By default, port 1701 is used for L2TP VPN-related communications. If required, active sessions can be stopped from this view. Use **FortiView > All Sessions**.

Testing L2TP VPN connections

To confirm that a VPN between a local network and a dialup client has been configured correctly, at the dialup client, issue a ping command to test the connection to the local network. The VPN tunnel initializes when the dialup client attempts to connect.

Logging L2TP VPN events

You can configure the FortiGate unit to log VPN events. For L2TP VPNs, connection events and tunnel status (up/down) are logged.

To log VPN events - GUI

1. Go to **Log & Report > Log Settings**.
2. Enable the storage of log messages to one or more locations.
3. Select **Enable**, and then select **VPN activity event**.
4. Select **Apply**.

To log VPN events - CLI

```
config log memory setting
    set diskfull overwrite
    set status enable
end
config log eventfilter
    set vpn enable
end
```

Session helpers

The FortiOS firewall can analyze most TCP/IP protocol traffic by comparing packet header information to security policies. This comparison determines whether to accept or deny the packet and the session that the packet belongs to.

Some protocols include information in the packet body (or payload) that must be analyzed to successfully process sessions for this protocol. For example, the SIP VoIP protocol uses TCP control packets with a standard destination port to set up SIP calls. But the packets that carry the actual conversation can use a variety of UDP protocols with a variety of source and destination port numbers. The information about the protocols and port numbers used for a SIP call is contained in the body of the SIP TCP control packets. To successfully process SIP VoIP calls, FortiOS must be able to extract information from the body of the SIP packet and use this information to allow the voice-carrying packets through the firewall.

FortiOS uses session helpers to analyze the data in the packet bodies of some protocols and make adjustments to allow those protocols to send packets through the firewall.

This section includes the topics:

- [Viewing the session helper configuration](#)
- [Changing the session helper configuration](#)
- [DCE-RPC session helper \(dcerpc\)](#)
- [DNS session helpers \(dns-tcp and dns-udp\)](#)
- [File transfer protocol \(FTP\) session helper \(ftp\)](#)
- [H.323 and RAS session helpers \(h323 and ras\)](#)
- [H.245 session helper \(h245\)](#)
- [Media Gateway Controller Protocol \(MGCP\) session helper \(mgcp\)](#)
- [ONC-RPC portmapper session helper \(pmap\)](#)
- [PPTP session helper for PPTP traffic \(pptp\)](#)
- [Remote shell session helper \(rsh\)](#)
- [Real-Time Streaming Protocol \(RTSP\) session helper \(rtsp\)](#)
- [Session Initiation Protocol \(SIP\) session helper \(sip\)](#)
- [Trivial File Transfer Protocol \(TFTP\) session helper \(tftp\)](#)
- [Oracle TNS listener session helper \(tns\)](#)

Viewing the session helper configuration

You can view the session helpers enabled on your FortiGate unit in the CLI using the commands below. The following output shows the first two session helpers. The number of session helpers can vary to around 20.

```
show system session-helper
config system session-helper
edit 1
    set name pptp
    set port 1723
```

```
    set protocol 6
  next
    set name h323
    set port 1720
    set protocol 6
  end
  .
  .
```

The configuration for each session helper includes the name of the session helper and the port and protocol number on which the session helper listens for sessions. Session helpers listed on protocol number 6 (TCP) or 17 (UDP). For a complete list of protocol numbers see [Assigned Internet Protocol Numbers](#).

For example, the output above shows that FortiOS listens for PPTP packets on TCP port 1723 and H.323 packets on port TCP port 1720.

If a session helper listens on more than one port or protocol the more than one entry for the session helper appears in the `config system session-helper` list. For example, the `pmap` session helper appears twice because it listens on TCP port 111 and UDP port 111. The `rsh` session helper appears twice because it listens on TCP ports 514 and 512.

Changing the session helper configuration

Normally you will not need to change the configuration of the session helpers. However in some cases you may need to change the protocol or port the session helper listens on.

Changing the protocol or port that a session helper listens on

Most session helpers are configured to listen for their sessions on the port and protocol that they typically use. If your FortiGate unit receives sessions that should be handled by a session helper on a non-standard port or protocol you can use the following procedure to change the port and protocol used by a session helper. The following example shows how to change the port that the `pmap` session helper listens on for Sun RPC portmapper TCP sessions. By default `pmap` listens on TCP port 111.

To change the port that the `pmap` session helper listens on to TCP port 112

1. Confirm that the TCP `pmap` session helper entry is 11 in the session-helper list:

```
show system session-helper 11
config system session-helper
  edit 11
    set name pmap
    set port 111
    set protocol 6
  next
end
```

2. Enter the following command to change the TCP port to 112.

```
config system session-helper
  edit 11
    set port 112
  end
```

3. The `pmap` session helper also listens on UDP port 111. Confirm that the UDP `pmap` session helper entry is 12 in the session-helper list:


```
show system session-helper 12
config system session-helper
edit 12
    set name pmap
    set port 111
    set protocol 17
next
end
```

4. Enter the following command to change the UDP port to 112.

```
config system session-helper
edit 12
    set port 112
end
```

Use the following command to set the h323 session helper to listen for ports on the UDP protocol.

To change the protocol that the h323 session helper listens on

1. Confirm that the h323 session helper entry is 2 in the session-helper list:

```
show system session-helper 2
config system session-helper
edit 2
    set name h323
    set port 1720
    set protocol 6
next
end
```

2. Enter the following command to change the protocol to UDP.

```
config system session-helper
edit 2
    set protocol 17
end
```

If a session helper listens on more than one port or protocol, then multiple entries for the session helper must be added to the session helper list, one for each port and protocol combination. For example, the rtsp session helper listens on TCP ports 554, 7070, and 8554 so there are three rtsp entries in the session-helper list. If your FortiGate unit receives rtsp packets on a different TCP port (for example, 6677) you can use the following command to configure the rtsp session helper to listen on TCP port 6677.

To configure a session helper to listen on a new port and protocol

```
config system session-helper
edit 0
    set name rtsp
    set port 6677
    set protocol 6
end
```

Disabling a session helper

In some cases you may need to disable a session helper. Disabling a session helper just means removing it from the session-helper list so that the session helper is not listening on a port. You can completely disable a session

helper by deleting all of its entries from the session helper list. If there are multiple entries for a session helper on the list you can delete one of the entries to prevent the session helper from listening on that port.

To disable the mgcp session helper from listening on UDP port 2427

1. Enter the following command to find the mgcp session helper entry that listens on UDP port 2427:

```
show system session-helper
.
.
.
edit 19
  set name mgcp
  set port 2427
  set protocol 17
next
.
.
.
```

2. Enter the following command to delete session-helper list entry number 19 to disable the mgcp session helper from listening on UDP port 2427:

```
config system session-helper
  delete 19
```

By default the mgcp session helper listens on UDP ports 2427 and 2727. The previous procedure shows how to disable the mgcp protocol from listening on port 2427. The following procedure completely disables the mgcp session helper by also disabling it from listening on UDP port 2727.

To completely disable the mgcp session helper

1. Enter the following command to find the mgcp session helper entry that listens on UDP port 2727:

```
show system session-helper
.
.
.
edit 20
  set name mgcp
  set port 2727
  set protocol 17
next
.
.
.
```

2. Enter the following command to delete session-helper list entry number 20 to disable the mgcp session helper from listening on UDP port 2727:

```
config system session-helper
  delete 20
```

DCE-RPC session helper (dcerpc)

Distributed Computing Environment Remote Procedure Call (DCE-RPC) provides a way for a program running on one host to call procedures in a program running on another host. DCE-RPC (also called MS RPC for Microsoft RPC) is similar to ONC-RPC. Because of the large number of RPC services, for example, MAPI, the transport

address of an RPC service is dynamically negotiated based on the service program's universal unique identifier (UUID). The Endpoint Mapper (EPM) binding protocol in FortiOS maps the specific UUID to a transport address.

To accept DCE-RPC sessions you must add a security policy with service set to any or to the DEC-RPC pre-defined service (which listens on TCP and UDP ports 135). The dcerpc session helper also listens on TCP and UDP ports 135.

The session allows FortiOS to handle DCE-RPC dynamic transport address negotiation and to ensure UUID-based security policy enforcement. You can define a security policy to permit all RPC requests or to permit by specific UUID number.

In addition, because a TCP segment in a DCE-RPC stream might be fragmented, it might not include an intact RPC PDU. This fragmentation occurs in the RPC layer; so FortiOS does not support parsing fragmented packets.



The DCE-RPC session helper does not support destination NAT (DNAT) or Firewall VIPs unless you are using the OXID Resolver service (also called IOXIDResolver).

DNS session helpers (dns-tcp and dns-udp)

FortiOS includes two DNS session helpers, dns-tcp, a session helper for DNS over TCP, and dns-udp, a session helper for DNS over UDP.

To accept DNS sessions you must add a security policy with service set to any or to the DNS pre-defined service (which listens on TCP and UDP ports 53). The dns-udp session helper also listens on UDP port 53. By default the dns-tcp session helper is disabled. If needed you can use the following command to enable the dns-tcp session helper to listen for DNS sessions on TCP port 53:

```
config system session-helper
  edit 0
    set name dns-tcp
    set port 53
    set protocol 6
  end
```

File transfer protocol (FTP) session helper (ftp)

The FTP session helper monitors PORT, PASV and 227 commands and NATs the IP addresses and port numbers in the body of the FTP packets and opens ports on the FortiGate unit as required.

To accept FTP sessions you must add a security policy with service set to any or to the FTP, FTP_Put, and FTP_GET pre-defined services (which all listen on TCP port 21).

H.323 and RAS session helpers (h323 and ras)

The H.323 session helper supports secure H.323 voice over IP (VoIP) sessions between terminal endpoints such as IP phones and multimedia devices. In H.323 VoIP networks, gatekeeper devices manage call registration, admission, and call status for VoIP calls. The FortiOS h323 session helper supports gatekeepers installed on two different networks or on the same network.

To accept H.323 sessions, you must add a security policy with service set to any or to the H323 pre-defined service (which listens on TCP port numbers 1720 and 1503 and on UDP port number 1719). The h323 session helper listens on TCP port 1720.

The ras session helper is used with the h323 session helper for H.323 Registration, Admission, and Status (RAS) services. The ras session helper listens on UDP port 1719.

Alternate H.323 gatekeepers

The h323 session helper supports using H.323 alternate gatekeepers. All the H.323 endpoints must register with a gatekeeper through the Registration, Admission, and Status (RAS) protocol before they make calls. During the registration process, the primary gatekeeper sends Gatekeeper Confirm (GCF) and Registration Confirm (RCF) messages to the H.323 endpoints that contain the list of available alternate gatekeepers.

The alternate gatekeeper provides redundancy and scalability for the H.323 endpoints. If the primary gatekeeper fails the H.323 endpoints registered with that gatekeeper are automatically registered with the alternate gatekeeper. To use the H.323 alternate gatekeeper, you need to configure security policies that allow H.323 endpoints to reach the alternate gatekeeper.

H.245 session helper (h245)

H.245 is a control channel protocol used for H.323 and other similar communication sessions. H.245 sessions transmit non-telephone signals. H.245 sessions carry information needed for multimedia communication, such as encryption, flow control jitter management and others.

FortiOS includes one H.245 session helper, h245, which serves call in and call out sessions. There is no standard port for H.245. The ports are negotiated through H.323 setup and connect messages. As administrator, you should configure the RAS or H.323 helper if non-standard ports are used for RAS or H323. You do not have to configure the h245 helper because the h323 helper configures the h245 helper.

Media Gateway Controller Protocol (MGCP) session helper (mgcp)

The Media Gateway Control Protocol (MGCP) is a text-based application layer protocol used for VoIP call setup and control. MGCP uses a master-slave call control architecture in which the media gateway controller uses a call agent to maintain call control intelligence, while the media gateways perform the instructions of the call agent.

To accept MGCP sessions you must add a security policy with service set to any or to the MGCP predefined service (which listens on UDP port numbers 2427 and 2727). The h323 session helper also listens on UDP port numbers 2427 and 2727.

The MGCP session helper does the following:

- VoIP signaling payload inspection. The payload of the incoming VoIP signaling packet is inspected and malformed packets are blocked.
- Signaling packet body inspection. The payload of the incoming MGCP signaling packet is inspected according to RFC 3435. Malformed packets are blocked.
- Stateful processing of MGCP sessions. State machines are invoked to process the parsed information. Any out-of-state or out-of-transaction packet is identified and properly handled.
- MGCP Network Address Translation (NAT). Embedded IP addresses and ports in packet bodies is properly translated based on current routing information and network topology, and is replaced with the translated IP address and port number, if necessary.

- Manages pinholes for VoIP traffic. To keep the VoIP network secure, the IP address and port information used for media or signaling is identified by the session helper, and pinholes are dynamically created and closed during call setup.

ONC-RPC portmapper session helper (pmap)

Open Network Computing Remote Procedure Call (ONC-RPC) is a widely deployed remote procedure call system. Also called Sun RPC, ONC-RPC allows a program running on one host to call a program running on another. The transport address of an ONC-RPC service is dynamically negotiated based on the service's program number and version number. Several binding protocols are defined for mapping the RPC program number and version number to a transport address.

To accept ONC-RPC sessions you must add a security policy with service set to any or to the ONC-RPC pre-defined service (which listens on TCP and UDP port number 111). The RPC portmapper session helper (called pmap) handles the dynamic transport address negotiation mechanisms of ONC-RPC.

PPTP session helper for PPTP traffic (pptp)

The PPTP session help supports port address translation (PAT) for PPTP traffic. PPTP provides IP security at the Network Layer. PPTP consists of a control session and a data tunnel. The control session runs over TCP and helps in establishing and disconnecting the data tunnel. The data tunnel handles encapsulated Point-to-Point Protocol (PPP) packets carried over IP.

To accept PPTP sessions that pass through the FortiGate unit you must add a security policy with service set to any or to the PPTP pre-defined service (which listens on IP port 47 and TCP port 1723). The pptp session helper listens on TCP port 1723.

PPTP uses TCP port 1723 for control sessions and Generic Routing Encapsulation (GRE) (IP protocol 47) for tunneling the encapsulated PPP data. The GRE traffic carries no port number, making it difficult to distinguish between two clients with the same public IP address. PPTP uses the source IP address and the Call ID field in the GRE header to identify a tunnel. When multiple clients sharing the same IP address establish tunnels with the same PPTP server, they may get the same Call ID. The call ID value can be translated in both the control message and the data traffic, but only when the client is in a private network and the server is in a public network.

PPTP clients can either directly connect to the Internet or dial into a network access server to reach the Internet. A FortiGate unit that protects PPTP clients can translate the clients' private IP addresses to a pool of public IP addresses using NAT port translation (NAT-PT). Because the GRE traffic carries no port number for address translation, the pptp session helper treats the Call ID field as a port number as a way of distinguishing multiple clients.

After the PPTP establishing a TCP connection with the PPTP server, the client sends a start control connection request message to establish a control connection. The server replies with a start control connection reply message. The client then sends a request to establish a call and sends an outgoing call request message. FortiOS assigns a Call ID (bytes 12-13 of the control message) that is unique to each PPTP tunnel. The server replies with an outgoing call reply message that carries its own Call ID in bytes 12-13 and the client's call ID in bytes 14-15. The pptp session helper parses the control connection messages for the Call ID to identify the call to which a specific PPP packet belongs. The session helper also identifies an outgoing call request message using the control message type field (bytes 8-9) with the value 7. When the session helper receives this message, it parses the control message for the call ID field (bytes 12-13). FortiOS translates the call ID so that it is unique across multiple calls from the same translated client IP. After receiving outgoing call response message, the session helper holds this message and opens a port that accepts GRE traffic that the PPTP server sends. An outgoing call request message contains the following parts:

- The protocol used for the outgoing call request message (usually GRE)
- Source IP address (PPTP server IP)
- Destination IP address (translated client IP)
- Destination port number (translated client call ID)

The session helper identifies an outgoing call reply message using the control message type field (bytes 8-9) with the value 8. The session helper parses these control messages for the call ID field (bytes 12-13) and the client's call ID (bytes 14-15). The session helper then uses the client's call ID value to find the mapping created for the other direction, and then opens a pinhole to accept the GRE traffic that the client sends.

An outgoing call reply message contains the following parts:

- Protocol used for the outgoing call reply message (usually GRE)
- Source IP address (PPTP client IP)
- Destination IP address (PPTP server IP)
- Destination port number (PPTP server Call ID)

Each port that the session opens creates a session for data traffic arriving in that direction. The session helper opens the following two data sessions for each tunnel:

- Traffic from the PPTP client to the server, using the server's call ID as the destination port
- Traffic from the PPTP server to the client, using the client's translated call ID as the destination port

The default timeout value of the control connection is 30 minutes. The session helper closes the pinhole when the data session exceeds the timeout value or is idle for an extended period.

Remote shell session helper (rsh)

Using the remote shell program (RSH), authenticated users can run shell commands on remote hosts. RSH sessions most often use TCP port 514. To accept RSH sessions you must add a security policy with service set to any or to the RSH pre-defined service (which listens on TCP port number 514).

FortiOS automatically invokes the rsh session helper to process all RSH sessions on TCP port 514. The rsh session helper opens ports required for the RSH service to operate through a FortiGate unit running NAT or transparent and supports port translation of RSH traffic.

Real-Time Streaming Protocol (RTSP) session helper (rtsp)

The Real-Time Streaming Protocol (RTSP) is an application layer protocol often used by SIP to control the delivery of multiple synchronized multimedia streams, for example, related audio and video streams. Although RTSP is capable of delivering the data streams itself it is usually used like a network remote control for multimedia servers. The protocol is intended for selecting delivery channels (like UDP, multicast UDP, and TCP) and for selecting a delivery mechanism based on the Real-Time Protocol (RTP). RTSP may also use the SIP Session Description Protocol (SDP) as a means of providing information to clients for aggregate control of a presentation consisting of streams from one or more servers, and non-aggregate control of a presentation consisting of multiple streams from a single server.

To accept RTSP sessions you must add a security policy with service set to any or to the RTSP pre-defined service (which listens on TCP ports 554, 770, and 8554 and on UDP port 554). The rtsp session helper listens on TCP ports 554, 770, and 8554.

The rtsp session help is required because RTSP uses dynamically assigned port numbers that are communicated in the packet body when end points establish a control connection. The session helper keeps track of the port

numbers and opens pinholes as required. In Network Address Translation (NAT) mode, the session helper translates IP addresses and port numbers as necessary.

In a typical RTSP session the client starts the session (for example, when the user selects the Play button on a media player application) and establishes a TCP connection to the RTSP server on port 554. The client then sends an OPTIONS message to find out what audio and video features the server supports. The server responds to the OPTIONS message by specifying the name and version of the server, and a session identifier, for example, 24256-1.

The client then sends the DESCRIBE message with the URL of the actual media file the client wants to play. The server responds to the DESCRIBE message with a description of the media in the form of SDP code. The client then sends the SETUP message, which specifies the transport mechanisms acceptable to the client for streamed media, for example RTP/RTCP or RDT, and the ports on which it receives the media.

In a NAT configuration the rtsp session helper keeps track of these ports and addresses translates them as necessary. The server responds to the SETUP message and selects one of the transport protocols. When both client and server agree on a mechanism for media transport the client sends the PLAY message, and the server begins streaming the media.

Session Initiation Protocol (SIP) session helper (sip)

The sip session helper is described in the VoIP Solutions: SIP Guide.

Trivial File Transfer Protocol (TFTP) session helper (tftp)

To accept TFTP sessions you must add a security policy with service set to any or to the TFTP pre-defined service (which listens on UDP port number 69). The TFTP session helper also listens on UTP port number 69.

TFTP initiates transfers on UDP port 69, but the actual data transfer ports are selected by the server and client during initialization of the connection. The tftp session helper reads the transfer ports selected by the TFTP client and server during negotiation and opens these ports on the firewall so that the TFTP data transfer can be completed. When the transfer is complete the tftp session helper closes the open ports.

Oracle TNS listener session helper (tns)

The Oracle Transparent Network Substrate (TNS) listener listens on port TCP port 1521 for network requests to be passed to a database instance. The Oracle TNS listener session helper (tns) listens for TNS sessions on TCP port 1521. TNS is a foundation technology built into the Oracle Net foundation layer and used by SQLNET.

Advanced concepts

This section provides configuration concepts and techniques to enhance your network security and includes the topics:

- [Single firewall vs. multiple virtual domains](#)
- [Modem](#)
- [FortiExtender](#)
- [Assigning IP address by MAC address](#)
- [IP addresses for self-originated traffic](#)
- [Disk](#)
- [CLI scripts](#)
- [Rejecting PING requests](#)
- [Opening TCP 113](#)
- [Obfuscate HTTP responses](#)

To see a collection of practical articles, go to Fortinet's [Cookbook](#) site and navigate to [Resources > SysAdmin Notes](#).

Single firewall vs. multiple virtual domains

A typical FortiGate setup, with a small to mid-range appliance, enables you to include a number of subnets on your network using the available ports and switch interfaces. This can potentially provide a means of having three or more mini networks for the various groups in a company. Within this infrastructure, multiple network administrators have access to the FortiGate to maintain security policies.


However, the FortiGate unit may not have enough interfaces to match the number of departments in the organization. If the FortiGate unit is running in transparent mode however, there is only one interface, and multiple network branches through the FortiGate are not possible.

A FortiGate unit with Virtual Domains (VDOMs) enabled, provides a means to provide the same functionality in transparent mode as a FortiGate in NAT mode. VDOMs are a method of dividing a FortiGate unit into two or more virtual units that function as multiple independent units. VDOMs can provide separate security policies and, in NAT mode, completely separate configurations for routing and VPN services for each connected network. For administration, an administrator can be assigned to each VDOM, minimizing the possibility of error or fouling network communications.


By default, most FortiGate units support 10 VDOMs. Many FortiGate models support purchasing a license key to increase the maximum number.

Single firewall vs. vdoms

When VDOMs are not enabled, and the FortiGate unit is in transparent mode, all the interfaces on your unit become broadcast interfaces. The problem is there are no interfaces free for additional network segments.



A FortiGate with three interfaces means only limited network segments are possible without purchasing more FortiGate devices.



With multiple VDOMs you can have one of them configured in transparent mode, and the rest in NAT mode. In this configuration, you have an available transparent mode FortiGate unit you can drop into your network for troubleshooting, and you also have the standard.

This example shows how to enable VDOMs on the FortiGate unit and the basic and create a VDOM accounting on the DMZ2 port and assign an administrator to maintain the VDOM. First enable Virtual Domains on the FortiGate unit.

To enable VDOMs - web-based manager

1. Go to **System > Dashboard > Status**.
2. In the **System Information** widget, select **Enable** for **Virtual Domain**.

Note that on FortiGate-60 series and lower models, you need to enable VDOMs in the CLI only.

The FortiGate unit logs you out. Once you log back in, you will notice that the menu structure has changed. This reflects the global settings for all Virtual Domains.

To enable VDOMs - CLI

```
config system global
    set vdom-admin enable
end
```

Next, add the VDOM called accounting.

To add a VDOM - web-based manager

1. Go to **Global > VDOM > VDOM**, and select **Create New**.
2. Enter the VDOM name `accounting`.
3. Select **OK**.

To add a VDOM - CLI

```
config vdom
  edit <new_vdom_name>
end
```

With the Virtual Domain created, you can assign a physical interface to it, and assign it an IP address.

To assign physical interface to the accounting Virtual Domain - web-based manager

1. Go to **Global > Network > Interface**.
2. Select the DMZ2 port row and select **Edit**.
3. For the **Virtual Domain** drop-down list, select **accounting**.
4. Select the **Addressing Mode** of **Manual**.
5. Enter the IP address for the port of 10.13.101.100/24.
6. Set the **Administrative Access** to **HTTPS** and **SSH**.
7. Select **OK**.

To assign physical interface to the accounting Virtual Domain - CLI

```
config global
  config system interface
    edit dmz2
      set vdom accounting
      set ip 10.13.101.100/24
      set allowaccess https ssh
    next
  end
```

Modem

FortiGate units support the use of wireless, 3G and 4G modems connected using the USB port or, if available, the express card slot. Modem access provides either primary or secondary (redundant) access to the Internet. For FortiGate units that do not include an internal modem (those units with an “M” designation), the modem interface will not appear in the web-based manager until enabled in the CLI. To enable the modem interface enter the CLI commands:

```
config system modem
  set status enable
end
```

You will need to log out of the FortiGate and log back in to see the modem configuration page at **Network > Modem**. Once enabled, modem options become available by going to **Network > Interfaces**.

Note that the modem interface is only available when the FortiGate unit is in NAT mode.

To configure modem settings, go to **> Network > Modem**.

Configuring the modem settings is a matter of entering the ISP phone number, user name and password. Depending on the modem, additional information may need to be supplied such as product identifiers, and initialization strings.

The FortiGate unit includes a number of common modems within its internal database. You can view these by selecting the **Configure Modem** link on the **Modem Settings** page. If your modem is not on the list, select **Create New** to add the information. This information is stored on the device, and will remain after a reboot.

Fortinet has an online database of modem models and configuration settings through FortiGuard. A subscription to the FortiGuard services is not required to access the information. As models are added, you can select the **Configure Modem** link and select **Update Now** to download new configurations.

USB modem port

Each USB modem has a specific dial-out port. This will be indicated with the documentation for your modem. To enable the correct USB port, use the CLI commands:

```
config system modem
    set wireless-port {0 | 1 | 2}
end
```

To test the port, use the diagnose command:

```
diagnose sys modem com /1
```

The 1 will be the value of your USB port selected. The response will be:

```
Serial port: /dev/l
Press Ctrl+W to exit.
```

If the port does not respond the output will be:

```
Can not open modem device '/dev/l' : Broken pipe
```

Modes

The FortiGate unit allows for two modes of operation for the modem; stand alone and redundant. In stand alone mode, the modem connects to a dialup ISP account to provide the connection to the Internet. In redundant mode, the modem acts as a backup method of connecting to the Internet, should the primary port for this function fails.

Configuring either stand alone or redundant modes are very similar. The primary difference is the selection of the interface that the modem will replace in the event of it failing, and the configuration of a PING server to monitor the chosen interface.

Configuring stand alone mode

Configuring stand alone mode is a matter of configuring the modem information and the dialing mode. The dial mode is either **Always Connect** or **Dial on demand**. Selecting **Always Connect** ensures that once the modem has connected, it remains connected to the ISP. Selecting **Dial on Demand**, the modem only calls the ISP if packets are routed to the modem interface. Once sent, the modem will disconnect after a specified amount of time.

To configure standalone mode as needed - GUI

1. Go to **Network > Modem**.
2. Select the **Mode** of **Standalone**.
3. Select the **Dial Mode** of **Dial on Demand**.
4. Select the number of redials the modem attempts if connection fails to 5.
5. Select **Apply**.

To configure standalone mode as needed- CLI

```
config system modem
    set status enable
    set mode standalone
    set auto-dial enable
    set redial 5
end
```

Configuring redundant mode

Redundant mode provides a backup to an interface, typically to the Internet. If that interface fails or disconnects, the modem automatically dials the configured phone number(s). Once connected, the FortiGate unit routes all traffic to the modem interface until the monitored interface is up again. The FortiGate unit pings the connection to determine when it is back online.

For the FortiGate to verify when the interface is back up, you need to configure a Ping server for that interface. You will also need to configure security policies between the modem interface and the other interfaces of the FortiGate unit to ensure traffic flow.

To configure redundant mode as needed - GUI

1. Go to **Network > Modem**.
2. Select the **Mode** of **Redundant**.
3. Select the interface the modem takes over from if it fails.
4. Select the **Dial Mode** of **Dial on Demand**.
5. Select the number of redials the modem attempts if connection fails to 5.
6. Select **Apply**.

To configure standalone mode as needed- CLI

```
config system modem
    set status enable
    set mode redundant
    set interface wan1
    set auto-dial enable
    set redial 5
end
```

Link Health Monitor

Adding a link health monitor is required for routing fail over traffic. A link health monitor will confirm the connectivity of the device's interface

To add a link health monitor

```
config system link-monitor
  edit "Example1"
    set srcint <Interface_sending_probe>
    set server <ISP_IP_address>
    set protocol <Ping or http>
    set gateway-ip <the_gateway_IP_to_reach_the_server_if_required>
    set failtime <failure_count>
    set interval <seconds>
    set update-cascade-interface enable
    set update-static-route enable
    set status enable
  end
```

Additional modem configuration

The CLI provides additional configuration options when setting up the modem options including adding multiple ISP dialing and initialization options and routing. For more information, see the CLI Reference.

Modem interface routing

The modem interface can be used in FortiOS as a dedicated interface. Once enabled and configured, you can use it in security policies and define static and dynamic routing. Within the CLI commands for the modem, you can configure the distance and priority of routes involving the modem interface. The CLI commands are:

```
config system modem
  set distance <route_distance>
  set priority <priority_value>
end
```

For more information on the routing configuration in the CLI, see the CLI Reference. For more information on routing and configuring routing, see the Advanced Routing Guide.

Assigning IP address by MAC address

To prevent users from changing their IP addresses and causing IP address conflicts or unauthorized use of IP addresses, you can bind an IP address to a specific MAC address using DHCP.

Use the CLI to reserve an IP address for a particular client identified by its device MAC address and type of connection. The DHCP server then always assigns the reserved IP address to the client. The number of reserved addresses that you can define ranges from 10 to 200 depending on the FortiGate model.

After setting up a DHCP server on an interface by going to **Network > Interfaces**, select the blue arrow next to **Advanced** to expand the options. If you know the MAC address of the system select **Create New** to add it, or if the system has already connected, locate it in the list, select the check box and select **Add from DHCP Client List**.

You can also match an address to a MAC address in the CLI. In the example below, the IP address 10.10.10.55 for User1 is assigned to MAC address 00:09:0F:30:CA:4F.

```
config system dhcp reserved-address
  edit User1
    set ip 10.10.10.55
    set mac 00:09:0F:30:CA:4F
    set type regular
  end
```

IP addresses for self-originated traffic

On the FortiGate unit, there are a number of protocols and traffic that is specific to the internal workings of FortiOS. For many of these traffic sources, you can identify a specific port/IP address for this self-originating traffic. The following traffic can be configured to a specific port/IP address:

- SNMP
- Syslog
- alert email
- FortiManager connection IP
- FortiGuard services
- FortiAnalyzer logging
- NTP
- DNS
- Authorization requests such as RADIUS
- FSSO

Configuration of these services is performed in the CLI. In each instance, there is a command `set source-ip`. For example, to set the source IP of NTP to be on the DMZ1 port with an IP of 192.168.4.5, the commands are:

```
config system ntp
  set ntpsync enable
  set syncinterval 5
  set source-ip 192.168.4.5
end
```

To see which services are configured with source-ip settings, use the `get` command:

```
get system source-ip status
```

The output will appear similar to the sample below:

```
NTP: x.x.x.x
DNS: x.x.x.x
SNMP: x.x.x.x
Central Management: x.x.x.x
FortiGuard Updates (AV/IPS): x.x.x.x
FortiGuard Queries (WebFilter/SpamFilter): x.x.x.x
```

Disk

To view the status and storage information of the local disk on your FortiGate unit, go to **System > Advanced**. The **DiskSettings** menu appears only on FortiGate units with an internal hard or flash disk.

Formatting the disk

The internal disk of the FortiGate unit (if available) can be formatted by going to **System > Advanced** and selecting **Disk Settings**.

Formatting the disk will erase all data on it, including databases for antivirus and IPS; logs, quarantine files, and WAN optimization caches. The FortiGate unit requires a reboot once the disk has been formatted.

Setting space quotas

If the FortiGate unit has an internal hard or flash disk, you can allocate the space on the disk for specific logging and archiving, and WAN optimization. By default, the space is used on an as required basis. As such, a disk can fill up with basic disk logging, leaving less potential space for quarantine.

By going to **System > Advanced**, you can select the **Edit** icon for **Logging and Archiving** and **WAN Optimization & Web Cache** and define the amount of space each log, archive and WAN optimization has on the disk.

CLI scripts

To upload bulk CLI commands and scripts, go to **System > Advanced**.

Scripts are text files containing CLI command sequences. Scripts can be used to deploy identical configurations to many devices. For example, if all of your devices use identical security policies, you can enter the commands required to create the security policies in a script, and then deploy the script to all the devices which should use those same settings.

Use a text editor such as Notepad or other application that creates simple text files. Enter the commands in sequence, with each line as one command, similar to examples throughout the FortiOS documentation set.

If you are using a FortiGate unit that is not remotely managed by a FortiManager unit or the FortiGuard Analysis and Management Service, the scripts you upload are executed and discarded. If you want to execute a script more than once, you must keep a copy on your management PC.

If your FortiGate unit is configured to use a FortiManager unit, you can upload your scripts to the FortiManager unit, and run them from any FortiGate unit configured to use the FortiManager unit. If you upload a script directly to a FortiGate unit, it is executed and discarded.

If your FortiGate unit is configured to use FortiGuard Analysis and Management Service, scripts you upload are executed and stored. You can run uploaded scripts from any FortiGate unit configured with your FortiGuard Analysis and Management Service account. The uploaded script files appear on the FortiGuard Analysis and Management Service portal web site.

Uploading script files

After you have created a script file, you can then upload it through **System > Advanced**. When a script is uploaded, it is automatically executed.

Commands that require the FortiGate unit to reboot when entered in the command line will also force a reboot if included in a script.

To execute a script

1. Go to **System > Advanced**.
2. Enable **Configuration Scripts**.
3. Select **Upload and Run a New Script** to locate the script file.
4. Select **Apply**.

If the FortiGate unit is not configured for remote management, or if it is configured to use a FortiManager unit, uploaded scripts are discarded after execution. Save script files to your management PC if you want to execute them again later.

If the FortiGate unit is configured to use the FortiGuard Analysis and Management Service, the script file is saved to the remote server for later reuse. You can view the script or run it from the FortiGuard Analysis and Management Service portal web site.

Auto repeat of CLI commands

Occasionally there is a need to repeatedly run a diagnose command over a long period of time (like checking CPU or memory usage, or checking proxy health). Previously, this could only be done with external console connections. With FortiOS 5.4.0, this can be done in a script using the `interval` and `repeat` commands.

Scripts can be uploaded as a file from the CLI or GUI. To upload scripts from the GUI go to **System > Advanced > Configuration Scripts** and upload and run the script.

To configure the schedule and scripts, use the following syntax:

```
config system auto-script
  edit <ScriptName>
    set interval
    set repeat
    set script
  end
end
```

`interval` the interval time in seconds between instances of the script running.

`repeat` the number of times to repeat the running of the script. The value 0 is used to set an infinite number of repetitions.

`start` **select** `manual` to start the script manually or `auto` to start the script automatically

`script` the contents of the script.

This feature may not be available on all models as a hard drive is necessary to make use of it.

CLI option to limit script output size

The CLI command `set output-size` limits the size of an auto script in megabytes and prevents the memory from being used up by the script's output.

CLI Syntax

```
config system auto-script
  edit <script name>
    set output-size <integer>
  next
end
```

Enter an integer value from 10 to 1024. Default is 10.

Execute restore script command

The `execute restore script` command merges arbitrary configlets into the running configuration from a script. You can carry out the command's authentication with either a username and password or with a certificate.

If the configuration in the script fails for any reason, the system will revert back to running configurations without interrupting the network.

To load script from the FTP, SCP, or TFTP servers, enter the following CLI command:

```
execute restore script {ftp | scp |  
tftp} <dir / filename in server> <server ip> <username> <password>
```

To view the results of the last restored script, enter the following CLI command:

```
execute restore script lastlog <dir / filename in server> <server ip> <username>  
<password>
```

Rejecting PING requests

The factory default configuration of your FortiGate unit allows the default external interface to respond to ping requests. Depending on the model of your FortiGate unit the actual name of this interface will vary. For the most secure operation, you should change the configuration of the external interface so that it does not respond to ping requests. Not responding to ping requests makes it more difficult for a potential attacker to detect your FortiGate unit from the Internet. One such potential threat are Denial of Service (DoS) attacks.

A FortiGate unit responds to ping requests if ping administrative access is enabled for that interface.

To disable ping administrative access - web-based manager

1. Go to **System > Network > Interface**.
2. Choose the external interface and select **Edit**.
3. Clear the **Ping Administrative Access** check box.
4. Select **OK**.

In the CLI, when setting the allowaccess settings, by selecting the access types and not including the PING option, that option is then not selected. In this example, only HTTPS is selected.

To disable ping administrative access - CLI

```
config system interface  
edit external  
set allowaccess https  
end
```

Opening TCP 113

Although seemingly contrary to conventional wisdom of closing ports from hackers, this port, which is used for ident requests, should be opened.

Port 113 initially was used as an authentication port, and later defined as an identification port (see RFC 1413). Some servers may still use this port to help in identifying users or other servers and establish a connection. Because port 113 receives a lot of unsolicited traffic, many routers, including on the FortiGate unit, close this port.

The issue arises in that unsolicited requests are stopped by the FortiGate unit, which will send a response saying that the port is closed. In doing so, it also lets the requesting server know there is a device at the given address, and thus announcing its presence. By enabling traffic on port 113, requests will travel to this port, and will most likely, be ignored and never responded to.

By default, the ident port is closed. To open it, use the following CLI commands:

```
config system interface
```

```
edit <port_name>
    set ident_accept enable
end
```

You could also further use port forwarding to send the traffic to a non-existent IP address and thus never have a response packet sent.

Obfuscate HTTP responses from SSL VPN

The FortiGate unit can obfuscate the HTTP responses from SSL VPN servers. By default this option is not enabled. To obfuscate HTTP headers, use the following CLI command:

```
config vpn ssl settings
    set url-obscurate {enable | disable}
end
```

Blocking land attacks in transparent mode

Enabling blocking land attacks allows BFD echo packets to pass through the FortiGate.

Since it's a system settings option you can enable or disable blocking land attacks for individual VDOMs if your FortiGate is operating with multiple VDOMs.

Another reason to enable this feature would be if your FortiGate is blocking BFD echo packets that should be allowed to pass through the FortiGate. For example, a FortiGate operating in transparent mode between two routers with a policy that allows all traffic may block BFD echo communication between the routers if blocking land attacks is disabled.

Use the following command to block land attacks and allow BFD echo packets. This option is disabled by default.

Syntax

```
config system settings
    set block-land-attack enable
end
```

Multi-dimension tagging

Multi-dimension tagging is available for address, device, and interface objects.

Tags can be configured in the GUI under **System > Tags**. Once created, these tags can be assigned to network interfaces, addresses, and custom devices and groups. In addition, a **Tags** column has been added for each of these GUI locations, including under **User & Device > Device Inventory**.

Syntax

The tagging option is available under `config user device` and `device-group`:

```
config user [device | device-group]
edit <name>
    config tagging
        edit <name>
            set category <object-tag-name>
            set tags <name>
        next
    end
```

```

    next
end

```

The same tagging option is also available under `config firewall multicast-address`, `multicast-address6`, `address`, `address6`, `addrgrp`, `addrgrp6`, `proxy-address`, and `config system interface` and `zone` (respectively):

```

config user [multicast-address | multicast-address6 | address | address6 | addrgrp |
  addrgrp6 | proxy-address]
  edit <name>
    config tagging
      edit <name>
        set category <object-tag-name>
        set tags <name>
      next
    end
  next
end

config system [interface | zone]
  edit <name>
    config tagging
      edit <name>
        set category <object-tag-name>
        set tags <name>
      next
    end
  next
end

```

Chapter 25 - Traffic Shaping

The following chapters are included in this document:

[The purpose of traffic shaping](#) describes traffic shaping theories and QoS.

[Traffic shaping methods](#) lists different methods of applying traffic shaping within FortiOS, and explains how to use Type of Service (ToS) and Differentiated Services (DiffServ).

[Examples](#) provides basic application scenarios for traffic shapers.

[Troubleshooting traffic shaping](#) lists diagnose commands that you can use to determine if traffic shapers are working correctly.

What's new in FortiOS 6.0

The following list contains new traffic shaping features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- ["Interface-based traffic shaping"](#) on page 2805
- ["Internet services support"](#) on page 2809

The purpose of traffic shaping

With the ever-increasing demands on network systems for a number of protocols, including email, HTTP traffic both internally and externally to the Internet, voice over IP, FTP, and more, slow traffic is becoming a reality. Important traffic may even be dropped or slowed to an unusable speed. Web traffic delays can result in a loss of revenue for businesses. Traffic shaping attempts to normalize traffic peaks and bursts to prioritize certain flows over others. There's a physical limitation to the amount of data that can be buffered and to the length of time it can be buffered.

FortiGate devices provide Quality of Service (QoS) by applying bandwidth limits and prioritization. You can use traffic shaping to adjust how a FortiGate allocates resources to different types of traffic to improve performance and stability of latency sensitive or bandwidth intensive network applications.

Traffic shaping, or traffic management, controls the bandwidth available and sets the priority of traffic processed by the policy to control the volume of traffic for a specific period (bandwidth throttling) or rate the traffic is sent (rate limiting).

Traffic shaping attempts to normalize traffic peaks and bursts to prioritize certain flows over others. But there is a physical limitation to the amount of data which can be buffered and to the length of time. Once these thresholds are surpassed, frames and packets are dropped, and sessions are affected in other ways.

A basic traffic shaping approach is to prioritize certain traffic flows over other traffic whose potential loss is less disadvantageous. This means that you accept certain sacrifices in performance and stability on low-priority traffic, to increase or guarantee performance and stability on high-priority traffic.

If, for example, you're applying bandwidth limitations to certain flows, you must accept the fact that these sessions can be limited and therefore negatively impacted.

Note that traffic shaping is effective for normal IP traffic at normal traffic rates. Traffic shaping isn't effective during periods when traffic exceeds the capacity of the FortiGate. Because packets must be received by the FortiGate before they're subject to traffic shaping, if the FortiGate can't process all of the traffic it receives, then dropped packets, delays, and latency are likely to occur.

To ensure that traffic shaping is working at its best, make sure that the interface Ethernet statistics show no errors, collisions, or buffer overruns.

Accelerated interfaces (NPx network processors and CE) affect traffic shaping. For more information, see the FortiOS [What's new in FortiOS 6.0.2](#) guide.

QoS

Quality of Service (QoS) is the capability to adjust some quality aspects of your overall network traffic. This can include such techniques as priority-based queuing and traffic policing. Because bandwidth is finite and because some types of traffic are slow, jitter or packet loss sensitive, bandwidth intensive, or operation critical, QoS can be a useful tool for optimizing the performance of various applications on your network.

Before implementing QoS, you should first identify the types of traffic that are important to your organization, the types of traffic that use high amounts of bandwidth, and the types of traffic that are sensitive to latency or packet loss.

For example, a company might want to guarantee sufficient bandwidth for revenue producing e-commerce traffic. They need to ensure that transactions can be completed and that clients don't experience service delays and

interruptions. At the same time, the company may need to ensure low latency for voice over IP (VoIP) traffic used by sales and customer support, while traffic latency and bursts may be less critical to the success of other network applications such as long term, resumable file transfers. Many organizations discover that QoS is especially important for managing their voice and streaming multimedia traffic. These types of traffic can rapidly consume bandwidth and are sensitive to latency.

Discovering the needs and relative importance of each traffic type on your network helps you to design an appropriate overall approach, including how you to configure each available QoS component technique. Some organizations discover that they only need to configure bandwidth limits for some services. Other organizations determine that they need to fully configure interface and security policy bandwidth limits for all services, and prioritize queuing of critical services relative to traffic rate.

You can implement QoS on FortiGate devices using the following techniques:

Technique	Description
Traffic policing	Drops packets that don't conform to bandwidth limitations
Traffic shaping	Ensures that the traffic may consume bandwidth at least at the guaranteed rate by assigning a greater priority queue if the guarantee isn't being met. Also ensures that the traffic can't consume bandwidth greater than the maximum at any given instance in time. Flows greater than the maximum rate are subject to traffic policing.
Queuing	Transmits packets in the order of their assigned priority queue for that physical interface. All traffic in a higher priority traffic queue must be completely transmitted before traffic in lower priority queues is transmitted.

When you're deciding how to configure QoS techniques, it can be helpful to know when FortiGate devices employ each technique in the overall traffic processing flow, and the considerations that arise from those mechanisms.

Traffic policing

A FortiGate begins to process traffic as it arrives (ingress) and departs (egress) on an interface. In later phases of network processing, such as enforcing maximum bandwidth use on sessions handled by a security policy, if the current rate for the destination interface or traffic regulated by that security policy is too high, the FortiGate may drop the packet. Time spent on prior processing, such as web filtering, decryption, or IPS, is often wasted on packets that aren't forwarded. This applies to VLAN interfaces and physical interfaces.

You can prevent this wasted effort on ingress by configuring the FortiGate to preemptively drop excess packets when they're received at the source interface, before most other traffic processing is performed:

```
config system interface
  edit <interface_name>
    set inbandwidth <limit>
  next
end
```

where <limit> is the bandwidth limit for incoming traffic, in kbps. Excess packets are dropped. If inbandwidth is 0, the rate isn't limited.

A similar CLI command is available that can be performed on egress:

```
config system interface
```

```
edit <interface_name>
    set outbandwidth <limit>
next
end
```

As with ingress, setting the rate to 0 (zero) sets the rate to unlimited.

Rate limiting traffic accepted by the interface allows you to restrict incoming traffic to rates that, while no longer the full capacity of the interface, at the traffic shaping point in the processing are more likely to result in acceptable rates of outgoing traffic per destination interface or all security policies. This conserves FortiGate processing resources for those packets that are more likely to be viable completely to the point of egress.

Excessive traffic policing can degrade network performance rather than improve it. For more information about factors that affect traffic policing, see [Important considerations on page 2789](#).

NP6 interfaces on FortiGate devices don't fully support bandwidth limits. When you set the outbandwidth setting on an NP6 interface, the FortiGate implements a lower bandwidth limit than the one that you configure. The inbandwidth setting has no effect on an NP6 interface, unless you disable NP offloading for the traffic on that interface.

Bandwidth guarantee, limit, and priority interactions

After packet acceptance, the FortiGate classifies traffic and may apply traffic policing at additional points during processing. It may also apply QoS techniques, such as prioritization and traffic shaping. Traffic shaping consists of a mixture of traffic policing to enforce bandwidth limits, and priority queue adjustment to assist packets in achieving the guaranteed rate.

If you have configured prioritization, the FortiGate prioritizes egressing packets by distributing them among FIFO (first in, first out) queues associated with each possible priority number. Each physical interface has six priority queues. Virtual interfaces don't have their own queues, and instead use the priority queues of the physical interface to which they are bound.

Each physical interface's six queues are queue 0 to queue 5, where queue 0 is the highest priority queue. However, for reasons described below, you may observe that your traffic uses only a subset of those six queues. Some traffic may always use a certain queue number. Some queuing may vary by the packet rate or mixture of services. Some queue numbers may be used only by through traffic for which you have configured traffic shaping in the security policy that applies to that traffic session. For example:

- Administrative access traffic will always use queue 0.
- Traffic matching security policies without traffic shaping may use queue 0, queue 1, or queue 2. Which queue will be used depends on the priority value you have configured for packets with that Type of Service (ToS) bit value, if you have configured ToS-based priorities.
- Traffic matching security policies with traffic shaping may use any queue. Which queue will be used depends on whether the packet rate is currently below the guaranteed bandwidth (queue 0), or above the guaranteed bandwidth. Packets at rates greater than the maximum bandwidth limit are dropped.
- If the global tos-based-priority is low (3), the priority in a traffic shaper is medium (2) and a packet flows through a policy that refers to the traffic shaper, the packet will be assigned the priority defined by the traffic shaper, in this case medium (2).

Prioritization and traffic shaping behavior varies according to your configuration, the service types and traffic volumes, and whether the traffic is through traffic, or the traffic originates from or terminates at the FortiGate itself.

FortiGate traffic

Security policies don't apply to administrative access to the FortiGate through HTTPS or SSH, or IPsec tunnel negotiations, and therefore FortiGate devices don't apply traffic shaping. Such traffic also uses the highest priority queue, queue 0. In other words, packet priority = 0.

Exceptions to this rule include traffic types that are connections related to a session governed by a security policy. For example, if you enabled scanning by FortiGuard antivirus, traffic from the sender technically terminates at the FortiGate proxy that scans that traffic type. The FortiGate initiates a second connection that transmits scanned content to its destination. Because the second connection's traffic is technically originating from the FortiGate proxy and therefore the FortiGate itself, it uses the highest priority queue, queue 0. However, this connection is logically associated with through traffic, and is therefore subject to possible bandwidth enforcement and guarantees in its governing security policy. In this way, it behaves partly like other through traffic.

Through traffic

For traffic passing through the FortiGate, the method a FortiGate uses to determine the priority queue varies by whether traffic shaping is enabled or not. Packets may or may not use a priority queue directly or indirectly derived from the Type of Service (ToS) bit — sometimes used instead with differentiated services — in the packet's IP header.

If traffic shaping isn't applied to a security policy, the FortiGate doesn't limit or guarantee bandwidth, and traffic for that session uses the priority queue determined directly by matching the ToS bit in its header with the configured values:

```
config system global
  set traffic-priority tos
  set traffic-priority-level {high | low | medium}
end
```

or, if you have configured a priority specifically for that ToS bit value:

```
config system tos-based-priority
  edit <id_int>
    set tos [0-15]
    set priority {high | low | medium}
  end
```

where `tos` is the value of the ToS bit in the packet's IP header, and `high` has a priority value of 0 and `low` is 2. Priority values configured in the second location will override the global ToS-based priority. In other words, packet priority = ToS-based priority.

For example, you might specify that packets with a ToS bit value of 2 should use queue 0, which is the highest priority queue:

```
config system tos-based-priority
  edit 15
    set tos 2
    set priority high
  next
end
```


If traffic shaping is applied to a security policy using a shared shaper, the FortiGate may subject packets to traffic policing or priority queue increases in an effort to meet bandwidth guarantees configured in the shaper.


For example, you might create a shared traffic shaper, where `high` has a priority value of 1 and `low` is 3, and `<rate>` is the bandwidth limit in kilobits per second:

```
config firewall shaper traffic-shaper
  edit <shaper_name>
    ...
    set priority {high | medium | low}
    set maximum-bandwidth <rate>
    set guaranteed-bandwidth <rate>
  end
```

Note that it's also necessary to create a traffic shaping policy and set it to use the shared traffic shaper:

```
config firewall shaping-policy
  edit <policy ID>
    ...
    set srcaddr <source address>
    set dstaddr <destination address>
    set service <service name>
    set dstintf <destination interface list>
    set traffic-shaper <shaper_name>
  end
```

The following diagram shows traffic queuing as the packet rate increases.



- If the current packet rate is less than guaranteed bandwidth, packets use priority queue 0: packet priority = 0.
- If the current packet rate is greater than guaranteed bandwidth but less than maximum bandwidth, the FortiGate assigns a priority queue by adding the numerical value of the security policy-based priority, where the value of `High` is 1, and `Low` is 3, with the numerical value of the ToS-based priority, where `high` has a priority value of 0 and `low` is 2. Because the two values are added, depending on the configured ToS-based priorities, packets in this category could use queues from queue 1 to queue 5. In other words: packet priority = ToS-based priority + security policy-based priority.

- If you enabled traffic shaping in the security policy, and the security policy's Traffic Priority is Low (value 3), and the priority normally applied to packets with that ToS bit is `medium` (value 1), then packets have a total packet priority of 4, and use priority queue 4.
- If the current packet rate exceeds the maximum bandwidth, excess packets are dropped.

Calculation and regulation of packet rates

Packet rates specified for Maximum Bandwidth or Guaranteed Bandwidth are:

$$\text{rate} = \text{amount} / \text{time}$$

where rate is expressed in Kbps

Burst size at any given instant can't exceed the amount configured for maximum bandwidth. Packets that exceed this are dropped. Packets deduct from the amount of bandwidth available to subsequent packets and available bandwidth regenerates at a fixed rate. As a result, bandwidth available to a given packet may be less than the configured rate, down to a minimum of 0 Kbps.


Rate calculation and behavior can alternatively be described using the token bucket metaphor, where:

- A traffic flow has an associated bucket, which represents burst size bounds, and is the size of your configured bandwidth limit.
- The bucket receives tokens, which represent available bandwidth, at the fixed configured rate.
- As time passes, tokens are added to the bucket, up to the capacity of the bucket. Excess tokens are discarded.
- When a packet arrives, the packet must deduct bandwidth tokens from the bucket equal to its packet size in order to egress.
- Packets can't egress if there are insufficient tokens to pay for its egress. These nonconforming packets are dropped.

Bursts aren't redistributed over a longer interval, so bursts are propagated rather than smoothed, although their peak size is limited.

Maximum burst size is the capacity of the bucket (the configured bandwidth limit). Actual size varies by the current number of tokens in the bucket, which may be less than bucket capacity, due to deductions from previous packets and the fixed rate at which tokens accumulate. A depleted bucket refills at the rate of your configured bandwidth limit. Bursts can't borrow tokens from other time intervals. This behavior is illustrated in the graph below.

Bursts and bandwidth limits over time



By limiting traffic peaks and token regeneration in this way, the available bandwidth at any given moment may be less than bucket capacity, but your limit on the total amount per time interval is ensured. Total bandwidth use during each interval of one second is, at most, the integral of your configured rate.

You may observe that external clients, such as FTP or BitTorrent clients, initially report rates between Maximum Bandwidth and twice that of Maximum Bandwidth, depending on the size of their initial burst. This is notably so when a connection is initiated following a period of no network activity. The apparent discrepancy in rates is caused by a difference in perspective when delimiting time intervals. A burst from the client may initially consume all tokens in the bucket, and before the end of one second, as the bucket regenerates, be allowed to consume almost another bucket's worth of bandwidth. From the perspective of the client, this constitutes one time interval. From the perspective of the FortiGate, however, the bucket can't accumulate tokens while full. Therefore, the time interval for token regeneration begins after the initial burst, and doesn't contain the burst. These different points of reference result in an initial discrepancy equal to the size of the burst — the client's rate contains it, but the FortiGate device's rate doesn't. If the connection is sustained to its limit and time progresses over an increasing number of intervals, however, this discrepancy decreases in importance relative to the bandwidth total, and the client's reported rate will eventually approach that of the configured rate limit for the FortiGate device.

For example, the Maximum Bandwidth might be 50 Kbps and there has been no network activity for one or more seconds. The bucket is full. A burst from an FTP client immediately consumes 50 kilobits. Because the bucket completely regenerates over one second, by the time almost another second has elapsed from the initial burst, traffic can consume another 49.999 kilobits, for a total of 99.999 kilobits between the two points in time. From the vantage point of an external FTP client regulated by this bandwidth limit, it therefore initially appears that the bandwidth limit is 99.999 Kbps, almost twice the configured limit of 50 Kbps. However, bucket capacity only regenerates at your configured rate of 50 Kbps, and so the connection can only consume a maximum of 50 kilobits during each second thereafter. The result is that as bandwidth consumption is averaged over an increasing number of time intervals, each of which are limited to 50 Kbps, the effects of the first interval's doubled

bandwidth size diminishes proportionately, and the client's reported rate eventually approaches your configured rate limit. The following table shows the effects of a 50 Kbps limit on client reported rates.

Total size transferred (kilobits)	Time (seconds)	Rate reported by client (Kbps)
99.999 (50 + 49.999)	1	99.999
149.999	2	74.999
199.999	3	66.666
249.999	4	62.499
299.999	5	59.998
349.999	6	58.333
...

Guaranteed Bandwidth can also be described using a token bucket metaphor. However, because this feature attempts to achieve or exceed a rate rather than limit it, the FortiGate doesn't discard non-conforming packets, as it does for Maximum Bandwidth. Instead, when the flow doesn't achieve the rate, the FortiGate increases the packets' priority queue, in an effort to increase the rate.

Guaranteed and maximum bandwidth rates apply to the bidirectional total for all sessions controlled by the security policy. For example, an FTP connection may entail two separate connections for the data and control portion of the session. Some packets may be reply traffic rather than initiating traffic. All packets for both connections are counted when calculating the packet rate for comparison with the guaranteed and maximum bandwidth rate.

Important considerations

By implementing Quality of Service (QoS), you trade some performance and/or stability from traffic X by discarding packets or introducing latency in order to improve performance and stability of traffic Y. The best traffic shaping configuration for your network balances the needs of each traffic flow by considering not only the needs of your particular organization, but also the resiliency and other characteristics of each particular service.

For example, you may find that web browsing traffic is both more resistant to interruptions or latency and less business critical than UDP or VoIP traffic, and so you might implement less restrictive QoS measures on UDP or VoIP traffic than on HTTP traffic.

An appropriate QoS configuration also takes into account the physical limits of your network devices and the interactions of the QoS mechanisms described in ["Bandwidth guarantee, limit, and priority interactions" on page 2784](#).

You may choose to configure QoS differently based on the hardware limits of your network and FortiGate. Traffic shaping may be less beneficial in extremely high-volume situations where traffic exceeds a network interface's or your FortiGate model's overall physical capacity. A FortiGate must have enough resources, such as memory and processing power, to process all traffic it receives, and to process it at the required rate. If it doesn't have this capacity, then dropped packets and increased latency are likely to occur. For example, if the total amount of memory available for queuing on a physical interface is frequently exceeded by your network's typical packet

rates, frames and packets must be dropped. In such a situation, you might choose to implement QoS using a higher model FortiGate, or to configure an incoming bandwidth limit on each interface.

Incorrect traffic shaping configurations can actually further degrade certain network flows, because excessive discarding of packets or increased latency beyond points that can be gracefully handled by that protocol can create additional overhead at upper layers of the network, which may be attempting to recover from these errors. For example, a configuration might be too restrictive on the bandwidth accepted by an interface, and may therefore drop too many packets, resulting in the inability to complete or maintain a SIP call.

To optimize traffic shaping performance, first ensure that the Ethernet statistics for the network interface are clean of errors, collisions, or buffer overruns. To check the interface, enter the following diagnose command to see the traffic statistics:

```
diagnose hardware deviceinfo nic <port_name>
```

If these aren't clean, adjust the FortiGate and settings of routers or other network devices that are connected to the FortiGate. For more information, see ["Troubleshooting traffic shaping" on page 2829](#).

Once Ethernet statistics are clean, you may want to use only some of the available FortiGate QoS techniques, or configure them differently, based on the nature of FortiGate QoS mechanisms described in ["Bandwidth guarantee, limit, and priority interactions" on page 2784](#).

Configuration considerations include:

- For maximum bandwidth limits, ensuring that bandwidth limits at the source interface and security policy aren't too low, which can cause the FortiGate to discard an excessive number of packets.
- For prioritization, considering the ratios of how packets are distributed between available queues, and which queue is used by which types of services. If you assign most packets to the same priority queue, it negates the effects of configuring prioritization. If you assign many high bandwidth services to high priority queues, lower priority queues may be starved for bandwidth and experience increased or indefinite latency. For example, you may want to prioritize a latency-sensitive service, such as SIP, over a bandwidth-intensive service such as FTP. Consider also that bandwidth guarantees can affect the queue distribution, assigning packets to queue 0 instead of their typical queue in high-volume situations.
- You may or may not want to guarantee bandwidth, because it causes the FortiGate to assign packets to queue 0 if the guaranteed packet rate isn't currently being met. Comparing queuing behavior for lower-bandwidth and higher-bandwidth situations, this would mean that effects of prioritization only become visible as traffic volumes rise and exceed their guarantees. Because of this, you might want only some services to use bandwidth guarantees, to avoid the possibility that in high-volume situations all traffic uses the same queue, thereby negating the effects of configuring prioritization.
- For prioritization, configure prioritization for all through traffic. You may want to configure prioritization by either ToS-based priority or security policy priority, but not both. This simplifies analysis and troubleshooting.

Traffic subject to both security policy and ToS-based priorities use a combined priority from both of those parts of the configuration, while traffic subject to only one of the prioritization methods will use only that priority. If you configure both methods, or if you configure either method for only a subset of your traffic, packets for which a combined priority applies will frequently receive a lower priority queue than packets for which you have only configured one priority method, or for which you have not configured prioritization.

For example, if both ToS-based priority and security policy priority both dictate that a packet should receive a medium priority, in the absence of bandwidth guarantees, a packet uses queue 3, while if only ToS-based priority is configured, the packet uses queue 1, and if only security policy-based priority is configured, the packet uses queue 2. If no prioritization is configured, the packet uses queue 0.

For example alternative QoS implementations that illustrate these considerations, see ["Examples" on page 2818](#).

Traffic shaping methods

There are three types of traffic shaping configurations in FortiOS. Each type has a specific function, and all types can be used together in varying configurations. Policy shaping allows you to define the maximum bandwidth and the guaranteed bandwidth set for a security policy. Per-IP traffic shaping allows you to define traffic control on a more granular level. Application traffic shaping goes further, allowing traffic controls on specific applications or application groupings.

This section describes the types of traffic shapers and how to configure them in the GUI and the CLI.



To configure traffic shaping in the GUI, you must enable **Traffic Shaping** in **System > Feature Visibility**.

Traffic shaping options

When you configure traffic shaping for your network, you can use the following methods to control the flow of network traffic to ensure that the traffic you want gets through, while also limiting bandwidth for less important traffic or traffic that consumes a lot of bandwidth.

- **Shared policy shaping** - bandwidth management by security policies
- **Per-IP shaping** - bandwidth management by user IP addresses
- **Application control shaping** - bandwidth management by application

Traffic shapers allow you to define how traffic will flow by setting the traffic priority, bandwidth, and DSCP options. You create shared policy traffic shapers and per-IP traffic shapers under Policy & Objects > Traffic Shapers.

You then enable traffic shapers within the traffic shaping policy, under Policy & Objects > Traffic Shaping Policy.

You can apply application control shaping to any traffic shaping policy, under Policy & Objects > Traffic Shaping Policy. You can control traffic by application category, application, and URL category.



To apply application control shaping, you must first enable application control at the policy level, under Policy & Objects > IPv4 Policy.

Traffic shaping policies allow you to apply traffic shaping measures to any traffic that matches your criteria. The criteria must specify a source, destination, service, and outgoing interface. Also, you must enable at least one type of traffic shaper to create a traffic shaping policy.

You can enable traffic shaping options on a FortiGate at the same time within a single traffic shaping policy. Generally, the hierarchy for traffic shapers in FortiOS is:

- Application control traffic shaper
- Shared policy traffic shaper
- Per-IP traffic shaper

Within this hierarchy, if an application control list has a traffic shaper defined, it has precedence over any other policy traffic shaper. For example, the Facebook application control example in [Application control shaping on page 2800](#) supersedes any security policy enabled traffic shapers. While the Facebook application may reach its maximum bandwidth, the user can still have the bandwidth room available from the shared traffic shaper and, if enabled, the per-IP traffic shaper.

Equally, any security policy shared traffic shaper has precedence over any per-IP traffic shaper. However, traffic that exceeds any of these traffic shapers is dropped. For example, the policy traffic shaper takes effect first, but if the per-IP traffic shaper limit is reached first, the traffic for that user is dropped even if the shared traffic shaper limit for the policy hasn't been exceeded.

Shared policy traffic shaping

Traffic shaping by security policy allows you to control the maximum and guaranteed throughput for any security policies specified in the traffic shaping policy.

When you configure a traffic shaper, you can apply bandwidth shaping per policy or for all policies. Depending on your selection, the FortiGate applies the traffic shaping rules differently.



By default, shared traffic shapers apply traffic shaping evenly to all policies that use. For **Per policy** and **All policies using this shaper** options to appear in the GUI, you must first enable it in the CLI. Go to Policy & Objects > Traffic Shapers and right-click on the traffic shaper to edit it in the CLI. Enter the following CLI commands:

```
set per-policy enable
end
```

Per policy

When you select a shared traffic shaper to be per policy, the FortiGate applies the traffic shaping rules to each security policy individually.

For example, if a traffic shaper is set to per policy with a maximum bandwidth of 1000 Kb/s and applied to four security policies, each policy has the same maximum bandwidth of 1000 Kb/s.

Per policy traffic shaping is compatible with client/server (active-passive) transparent mode WAN optimization rules. Traffic shaping is ignored for peer-to-peer WAN optimization and for client/server WAN optimization not operating in transparent mode.

For all policies using a traffic shaper

When you select a shared shaper to apply to all policies - **All Policies using this shaper** - the FortiGate applies the traffic shaping rules to all policies using the same shaper. For example, a traffic shaper is set to be per policy with a maximum bandwidth of 1000 Kbps. There are four security policies monitoring traffic through the FortiGate. All four have the traffic shaper enabled. Each security policy must share the defined 1000 Kbps, and is set on a first come, first served basis. For example, if policy 1 uses 800 Kbps, the remaining three must share 200 Kbps. As policy 1 uses less bandwidth, it's opened up to the other policies to use as required. Once used, any other policies encounters latency until free bandwidth opens from a policy currently in use.

Maximum and guaranteed bandwidth

The maximum bandwidth instructs the security policy what the largest amount of traffic allowed using the policy. Depending on the service or the users included for the security policy, this number can provide a larger or smaller throughput depending on the priority you set for the traffic shaper.

The **Maximum Bandwidth** can be set to a value of between 1 and 16776000 Kbps. The GUI gives an error if any value outside of this range is used, but in the CLI a value of 0 can be entered. Setting `maximum-bandwidth` to 0 (zero) prevents any traffic from going through the policy.

The guaranteed bandwidth ensures there's a consistent reserved bandwidth available for a given service or user. When setting the guaranteed bandwidth, ensure that the value is significantly less than the bandwidth capacity of the interface, otherwise no other traffic will pass through the interface or very little and potentially causing unwanted latency.

Traffic priority

Select a traffic priority of high, medium, or low, so the FortiGate manages the relative priorities of different types of traffic. For example, a policy for connecting to a secure web server that needs to support e-commerce traffic should be assigned a high traffic priority. Less important services should be assigned a low priority. The firewall provides bandwidth to low-priority connections only when bandwidth isn't needed for high-priority connections.

Be sure to enable traffic shaping on all security policies. If you don't apply a traffic shaping rule to a policy, the policy is set to high priority by default. Distribute security policies over all three priority queues.

Traffic shaping policy order

You must also place the traffic shaping policies in the correct order in the traffic shaping policy list page to get the desired results. It's necessary to arrange your traffic shaping policies into a sequence that places your more granular policies above general Internet access policies. For example, you should place any policies with application control shaping at the top of the traffic shaping policy list, followed by more general traffic shaping policies with shared policy shapers and per-IP traffic shapers.

The policy list page is located under Policy & Objects > Traffic Shaping Policy. To change the order of the policies, select the far left column to move the policy up or down. Make sure that the **Seq.#** column is showing on your menu so you can easily verify a policy's position in the sequence.

The following example shows how to order your policies. The high priority VoIP traffic shaping policy is placed at the top of the list, followed by restrictive policies to control streaming media, and your general Internet access policy last.

Seq.#	Source Address	Destination	Outgoing Interface	Shared Shaper	Reverse Shaper	Service	Application	Application Category
IPv4 (1 - 4)								
1	• all	• all	• wan1	guarantee-100kbps	guarantee-100kbps	• ALL	• SIP	• VoIP
2	• all	• all	• wan1	limited_bandwidth	limited_bandwidth	• ALL	• Vimeo • YouTube	• Video/Audio
3	• all	• all	• wan1	high-priority	high-priority	• ALL		

Traffic shaping policy configuration settings

To configure a traffic shaping policy, go to Policy & Objects > Traffic Shaping Policy and select the **Create New +** sign to create a new traffic shaping policy.

Set the **Matching Criteria** to the default options shown below or specify the criteria so that it matches a specific security policy.

Source	*all (default)
Destination	*all (default)
Service	*ALL (default)
Application Category	- Choose an application category to apply traffic shaping to a specific category of applications. For example, P2P, Social.Media, or VoIP.
Application	- Choose an application to specify which applications you want to apply traffic shaping to. For example, YouTube, Vimeo, or Facebook.
URL Category	- Choose a URL category to block a subset of applications. For example, potentially liable websites, security risks, or bandwidth consuming services.

Set **Apply shaper** to the following:

Outgoing Interface	*any (Set this to the external interface that you want to apply traffic shaping to. For example, wan1 is often used.)
Shared Shaper	Choose one of the default shared traffic shapers: guarantee-100kbps, high-priority, medium-priority, low-priority, shared-1M-pipe or create your own under Policy & Objects > Traffic Shapers. Shared traffic shapers share the allotted bandwidth with any security policies using them (unless they're set to per-policy in the CLI). This affects uploads or outbound traffic.
Reverse Shaper	Choose one of the default shared traffic shapers: guarantee-100kbps, high-priority, medium-priority, low-priority, shared-1M-pipe, or create your own under Policy & Objects > Traffic Shapers. This affects downloads or inbound traffic.
Per-IP Shaper	Enable a per-IP traffic shaper if you want to apply traffic shaping by bandwidth management by user IP addresses. You create traffic shapers under Policy & Objects > Traffic Shapers. Per-IP shapers affect downloads and uploads.
Enable this policy	Policies are enabled by default, but if you want to disable a traffic shaping policy de-select it here.

To create the traffic shaping policy - CLI:

```
config firewall shaping-policy
edit <shaping_policy_ID>
```

```

set srcaddr <source_address>
set dstaddr <destination_address>
set service <service_name>
set schedule {always | none}
application <application_name>
app-category <application_category_ID_list>
url-category <URL_category_ID_list>
dstintf <destination_interface_list>
traffic-shaper <shared_shaper_name>
traffic-shaper-reverse <reverse_traffic_shaper_name>
per-ip-shaper <per_IP_shaper_name>
end

```

VLAN, VDOM, and virtual interfaces

Policy-based traffic shaping doesn't use queues directly. It shapes the traffic and if the packet is allowed by the security policy, a priority is assigned. That priority controls what queue the packet is put in upon egress. VLANs, VDOMs, aggregate ports, and other virtual devices don't have queues and, as such, traffic is sent directly to the underlying physical device where it's queued and affected by the physical ports. This is also the case with IPsec connections.

Shared traffic shaper configuration settings

To configure a shared traffic shaper go to **Policy & Objects > Traffic Shapers** and select the **Create New +** sign to create a new traffic shaper.

Type	Select Shared .
Name	Enter a name for the traffic shaper.
Apply Shaper	<p>When selecting a traffic shaper to be Per Policy, the FortiGate applies the traffic shaping rules defined to each security policy individually. For example, if a traffic shaper is set to per policy, with a maximum bandwidth of 1000 Kbps, any security policies that have that traffic shaper enabled get 1000 Kbps of bandwidth each.</p> <p>When selecting a traffic shaper to apply to all policies (For All Policies Using This Shaper), the FortiGate applies the traffic shaping rules to all policies using the same traffic shaper. For example, the traffic shaper is set to be per policy with a maximum bandwidth of 1000 Kbps. There are four security policies monitoring traffic through the FortiGate. All four have the traffic shaper enabled. Each security policy must share the defined 1000 Kbps, and is set on a first come, first served basis. For example, if policy 1 uses 800 Kbps, the remaining three must share 200 Kbps. As policy 1 uses less bandwidth, that open bandwidth becomes available to the other policies to use as required. Once used, any other policies encounter latency until free bandwidth opens from a policy currently in use.</p>

Traffic Priority	<p>Select level of importance Priority so the FortiGate manages the relative priorities of different types of traffic. For example, a policy for connecting to a secure web server needed to support e-commerce traffic should be assigned a high traffic priority. Less important services should be assigned a low priority.</p> <p>If you don't apply any traffic shaping priority, the priority is set to high priority, by default.</p>
Maximum Bandwidth	<p>The maximum bandwidth instructs the security policy what the largest amount of traffic allowed using the policy. Depending on the service or the users included for the security policy, this number provides a larger or smaller throughput depending on the priority you set for the traffic shaper.</p> <p>Setting Maximum Bandwidth to 0 provides unlimited bandwidth.</p>
Guaranteed Bandwidth	<p>The guaranteed bandwidth ensures that a consistent reserved bandwidth is available for a given service or user. Ensure that you set the bandwidth to a value that's significantly less than the bandwidth capacity of the interface. Otherwise, little to no traffic passes through the interface and potentially causes unwanted latency.</p> <p>Setting Guaranteed Bandwidth to 0 provides unlimited bandwidth.</p>
DSCP	<p>Enter the number for the DSCP value. You can use the FortiGate Differentiated Services feature to change the DSCP (Differentiated Services Code Point) value for all packets accepted by a policy. The network can use these DSCP values to classify, mark, shape, and police traffic, and to perform intelligent queuing. DSCP features are applied to traffic by configuring the routers on your network to apply different service levels to packets depending on the DSCP value of the packet. For more information, see Traffic shaping methods.</p>

Shared traffic shaper per policy example

The following steps create a per policy traffic shaper called "Throughput" with a maximum traffic amount of 720,000 Kbps, and a guaranteed traffic of 150,000 Kbps with a high traffic priority.

To create the shared traffic shaper - GUI:

1. Go to **Policy & Objects > Traffic Shapers** and select the **Create New** + sign.
2. Set the **Type** to **Shared**.
3. Enter the **Name** *Throughput*.
4. Set the **Apply shaper** field to **Per Policy**.



By default, shared traffic shapers apply traffic shaping evenly to all policies that use it. For **Per policy** and **All policies using this shaper** options to appear in the GUI, you must first enable it in the CLI. Go to Policy & Objects > Traffic Shapers and right-click on the traffic shaper to edit it in the CLI. Enter the following CLI commands:

```
set per-policy enable
end
```

5. Set the **Traffic Priority** to **High**.
6. Select the **Maximum Bandwidth** check box and enter the value 150000.
7. Select the **Guaranteed Bandwidth** check box and enter the value 120000.
8. Select **OK**.

To create the shared shaper - CLI:

```
config firewall shaper traffic-shaper
edit Throughput
set per-policy enable
set maximum-bandwidth 150000
set guaranteed-bandwidth 120000
set priority high
end
```

Per-IP traffic shaping

Traffic shaping by IP allows you to apply traffic shaping to all source IP addresses in the security policy. In addition to controlling the maximum bandwidth users of a selected policy, you can also define the maximum number of concurrent sessions.

Per-IP traffic shaping allows you to limit the behavior of every member of a policy to avoid having one user use all of the available bandwidth. The bandwidth is shared equally within a group. Using a per-IP traffic shaper avoids having to create multiple policies for every user you want to apply a traffic shaper. Per-IP traffic shaping isn't supported over NP2 interfaces.

Per-IP traffic shaping configuration settings

To configure per-IP traffic shaping go to **Policy & Objects > Traffic Shapers > Per-IP** and select the **Create New** + sign.

Type	Select Per-IP .
Name	Enter a name for the per-IP traffic shaper.
Maximum Bandwidth	<p>The maximum bandwidth instructs the security policy what the largest amount of traffic allowed using the policy. Depending on the service or the users included for the security policy, this number can provide a larger or smaller throughput depending on the priority you set for the traffic shaper.</p> <p>Setting Maximum Bandwidth to 0 (zero) provides unlimited bandwidth.</p>

Maximum Concurrent Connections	Enter the maximum allowed concurrent connections.
Forward DSCP Reverse DSCP	Enter the number for the DSCP value. You can use the FortiGate Differentiated Services feature to change the DSCP (Differentiated Services Code Point) value for all packets accepted by a policy. The network can use these DSCP values to classify, mark, shape, and police traffic, and to perform intelligent queuing. DSCP features are applied to traffic by configuring the routers on your network to apply different service levels to packets depending on the DSCP value of the packet. For more information, see Traffic shaping methods .

Example

The following steps create a per-IP traffic shaper called “Accounting” with a maximum traffic amount of 720,000 Kbps, and the number of concurrent sessions of 200.

To create the shared shaper - GUI:

1. Go to **Policy & Objects > Traffic Shapers** and select the **Create New** “Plus” Icon.
2. Set the **Type** to **Per-IP**.
3. Enter the **Name** *Accounting*.
4. Enable the **Maximum Bandwidth** and enter the value *720000*.
5. Enable the **Maximum Concurrent Sessions** and enter the value *200*.
6. Select **OK**.

To create the shared shaper - CLI:

```
config firewall shaper per-ip-shaper
edit Accounting
set max100-bandwidth 720000
set max-concurrent-session 200
end
```

Adding a per-IP traffic shaper to a traffic shaping policy

Per-IP traffic shaping is supported by IPv6 security policies. You can add any per-IP traffic shaper to an IPv6 security policy in the CLI.

Example

The following steps show you how to add an existing per-IP traffic shaper to an IPv6 security policy. Make sure that you have already created a per-IP traffic shaper under Policy & Objects > Traffic Shapers.

To add a per-IP traffic shaper to an IPv6 security policy - GUI:

1. Go to **Policy & Objects > IPv6 Policy** and click the **Create New** + icon to create an Internet access policy.
2. Set the following:

Name	Enter a descriptive name
Incoming Interface	Internal
Source address	All
Outgoing interface	wan1
Destination address	all
Schedule	Always
Service	Any
Action	Accept

3. Select **OK**.
4. Go to Policy & Objects > Traffic Shaping Policy and the **Create New** + icon to create a new traffic shaping policy.
5. To apply your traffic shaping policy to the security policy you created earlier set the **Matching Criteria** to the following:

Source	all
Destination address	all
Service	ALL
Application Category	-
Application	-
URL Category	-

6. Under **Apply shaper**, set the following:

Outgoing interface	any (The outgoing interface should match the outgoing interface of the security policy you want to apply traffic shaping to.)
Shared Shaper	-
Reverse Shaper	-
Per-IP Shaper	Enable Per-IP Shaper and select your traffic traffic shaper from the drop-down menu.
Enable this policy	Enable this policy

7. Select **OK**.
8. On the policy list page, move the per-IP traffic shaper to the top of the list by clicking on the far left column to drag and drop it.

There are two methods to configure traffic shaping in the CLI. You can add a per-IP traffic shaper directly to an IPv6 security policy, or you can add a per-IP shaper to a traffic shaping policy. The second method will allow you to apply traffic shaping based on the interface and can therefore affect multiple security policies easily. The first method requires that you enable traffic shaping individually in all policies using the same two interfaces.

To add a per-IP traffic shaper to an IPv6 security policy- CLI:

```
config firewall policy6
  edit <security_policy_ID_number>
    set per-ip-shaper <per_IP_shaper_name>
  end
```

To add a per-IP traffic shaper to an IPv6 traffic shaping policy - CLI:

```
config firewall shaping-policy
  edit 1 <security_policy_ID_number>
    set ip-version 6
    set srcaddr <source_address>
    set dstaddr <destination_address>
    set service <service_name>
    set dstintf <outgoing_interface>
    set per-ip-shaper <per_IP_shaper_name>
  end
```

Application control shaping

Traffic shaping is also possible for specific applications. Application control shaping works in conjunction with a shared traffic shaper or per-IP traffic shaper. You must create a traffic shaper with the bandwidth settings you would like to enforce or edit one of the predefined traffic shapers in the Policy & Objects > Traffic Shapers menu.

Traffic shaping policies allow you to enable these traffic shapers and configure application control options. In the traffic shaping policy, you can set an **Application Category**, **Application**, and **URL Category**. You must also specify which security policies to apply your traffic shaper to by setting the **Matching Criteria**. You can create a traffic shaping policy in the Policy & Objects > Traffic Shaping Policy section.



For application control shaping to work, application control must be enabled in a security policy, through Policy & Objects > IPv4 Policy or Policy & Objects > IPv6 Policy under **Security Profiles**.

Also, application control traffic shaping affects only applications that are set to pass in the Security Profiles > Application Control menu.

For more information about application control, see the FortiOS [Chapter 21 - Security Profiles](#) Guide.

Example

This example sets the traffic shaping definition for Facebook to a medium priority, a default traffic shaper.

To add traffic shaping for Facebook - GUI:

1. Go to Policy & Objects > IPv4 Policy to create a general Internet access security policy.
2. Select the **Create New** + icon in the upper right corner of the screen to create a new security policy (or edit an

existing Internet access policy).

- Set the following to enable application control within a security policy:

Name	<Enter a descriptive name.>
Incoming Interface	Internal
Source address	All
Outgoing interface	wan1
Destination address	all
Schedule	Always
Service	Any
Action	Accept
Application Control	Under Security Profiles, enable Application Control and select the default application control profile.

- Select **OK**.
- Go to Policy & Objects > Traffic Shaping Policy and the **Create New** + icon to create a new traffic shaping policy.
- To apply your traffic shaping policy to the security policy you created earlier, set the **Matching Criteria** to the following:

Source	all
Destination address	all
Service	ALL
Application Category	Social.Media
Application	Facebook
URL Category	Social Networking

- Under **Apply shaper**, set the following:

Outgoing interface	any (The outgoing interface should match the outgoing interface of the security policy you want to apply shaping to.)
Shared Shaper	Enable Shared Shaper and select medium-priority from the drop-down menu.
Reverse Shaper	Enable Shared Shaper and select medium-priority from the drop-down menu.
Enable this policy	Enable this policy.

8. Select **OK**.
9. On the policy list page, move the Facebook traffic shaping policy to the top of the list by clicking on the far left column to drag and drop it.

To create a traffic shaping policy for Facebook - CLI:

```
config firewall shaping-policy
edit 1 <shaping_policy_ID_number>
set srcaddr all
set dstaddr all
set service ALL
set application 15832
set app-category 23 <Social.Media>
set url-category 37 <Social Networking>
set dstintf wan1 <outgoing_interface>
set traffic-shaper medium-priority
set reverse-traffic-shaper medium-priority
end
```

Reverse direction traffic shaping

The traffic shaper you select in the traffic shaping policy (shared traffic shaper) affects the traffic in the direction defined in the policy. For example, if the source port is lan and the destination is wan1, the traffic shaping affects the flow in this direction only — affecting the upload speed of the outbound traffic. By selecting **Shared Traffic Shaper Reverse Direction**, you can define the traffic shaper for the policy in the opposite direction to affect the download speed of the inbound traffic. In this example, from wan 1 to lan.

To add a reverse shaper - GUI:

1. Go to Policy & Objects > Traffic Shaping Policy.
2. Click **Create New** or select an existing policy and click **Edit**.
3. Set the **Matching Criteria** to match the interfaces of any security policies you want to affect.
4. Navigate to the **Apply shaper** section, enable the **Shared Shaper**, and select a traffic shaper from the drop-down menu.
5. Enable the **Reverse Shaper** and select a traffic shaper from the drop-down menu.
6. Select **OK**.

Setting the reverse direction only

There may be instances where you only need traffic shaping for incoming connections, which is in the reverse direction of typical traffic shapers.

To add a reverse traffic shaper - GUI:

1. Go to Policy & Objects > Traffic Shaping Policy.
2. Click **Create New** or select an existing policy and click **Edit**.
3. Set the **Matching Criteria** to match the interfaces of any security policies you want to affect.
4. Navigate to the **Apply shaper** section, enable the **Reverse Shaper** and select a traffic shaper from the drop-down menu.
5. Select **OK**.

To configure a reverse-only traffic shaper in a traffic shaping policy - CLI:

```
config firewall shaping-policy
  edit <policy_number>
    set reverse-traffic-shaper medium-priority
  end
```

To configure a reverse-only shaper within a security policy - CLI:

```
config firewall policy
  edit <policy_number>
    ...
    set traffic-shaper-reverse <shaper_name>
  end
```

Enabling traffic shaping in the security policy

Historically, FortiOS traffic shapers have always been enabled within a security policy. This is no longer the easiest way to apply traffic shapers, since in FortiOS 5.4 traffic shaping is now configured in the traffic shaping policy section, under Policy & Objects > Traffic Shaping Policy. However, you can still enable traffic shapers within a security policy using CLI commands and it will then appear in the GUI afterwards. The traffic shapers always go into effect after any DoS detection policies, and before any routing or packet scanning occurs.

Traffic shaping is also supported for IPv6 policies.



This isn't the recommended method, as it's easier to keep track of and order your traffic shaping policies if you configure them within a traffic shaping policy.

To enable traffic shaping within a security policy - CLI:

```
config firewall policy
  edit <policy_number>
    ...
    set traffic-shaper <traffic_shaper_name>
    set reverse-traffic-shaper <traffic_shaper_name>
    set per-ip-shaper <per_IP_traffic_shaper_name>
  end
```

Shared traffic shapers affect outbound traffic heading to a destination. To affect inbound traffic, or downloads, enable the **Reverse Shaper** also. For more information, see ["Reverse direction traffic shaping" on page 2802](#).

Scheduling traffic shaping policies

In FortiOS 5.6.3, a new scheduling feature was added to traffic shaping policies to apply different traffic shaping profiles at different times. This "schedule" attribute is currently only available using the CLI, and you can use this feature to apply a recurring schedule to your traffic shaping policies. The default recurring schedule options available are **always** or **none**. You can also create new schedules or schedule groups under **Policy & Objects > Schedules**. This allows you to create custom recurring or one-time schedules that can then be applied to your traffic shaping policies using the CLI commands below.

To schedule traffic shaping policies - CLI:

```
config firewall shaping-policy
edit <shaping_policy_ID>
set schedule {always | none}
end
```

ToS priority

Type of service (ToS) is an 8-bit field in the IP header that allows you to determine how the IP datagram should be delivered, using criteria of delay, throughput, priority, reliability, and cost. Each quality helps gateways determine the best way to route datagrams. A router maintains a ToS value for each route in its routing table. The lowest priority ToS is 0, and the highest is 7 when bits 3, 4, and 5 are all set to 1. There are other seldom used or reserved bits that aren't listed here.

Together these bits are the ToS variable of the `tos-based-priority` command. The router tries to match the ToS of the datagram to the ToS on one of the possible routes to the destination. If there's no match, the datagram is sent over a zero ToS route. Using increased quality may increase the cost of delivery because better performance may consume limited network resources.

Each bit represents the priority as defined in [RFC 1349](#):

- 1000 - minimize delay
- 0100 - maximize throughput
- 0010 - maximize reliability
- 0001 - minimize monetary cost

To set the ToS value - CLI:

```
config system tos-based-priority
edit <sequence_number>
set tos [0-15]
set priority {high | medium | low}
end
```

Where `tos` is the value of the type of service bit in the IP datagram header with a value between 0 and 15, and `priority` is the priority of this type of service priority. These priority levels conform to the firewall traffic shaping priorities, as defined in [RFC 1349](#).

For example, if you want to configure a FortiGate so that reliability is the first priority, set the ToS value to 4.

```
config system tos-based-priority
edit 1
set tos 4
set priority high
end
```

For a list of ToS values and their DSCP equivalents see [Traffic shaping methods on page 2791](#).

Example

```
config system tos-based-priority
edit 1
set tos 1
```

```
        set priority low
    next
    edit 4
        set tos 4
        set priority medium
    next
    edit 6
        set tos 6
        set priority high
    next
end
```

ToS in FortiOS

Traffic shaping and ToS follow the following sequence:

- The CLI command `tos-based-priority` acts as a `tos-to-priority` mapping. FortiOS maps the ToS to a priority when it receives a packet.
- Traffic shaping settings adjust a packet's priority according to the traffic.
- Deliver the packet based on its priority.

Traffic shaping units of measurement

Bandwidth speeds are measured in Kilobits per second (Kbps), and Bytes that are sent and received are measured in megabytes (MB). Occasionally this can cause confusion depending on whether your ISP uses kilobits per second (Kbps), kilobytes per second (KBps), megabits per second (Mbps), or gigabits per second (Gbps).

Download speeds

- 1 kilobyte per second (KBps) = 8 kilobits per second (Kbps)
- 1 megabit per second (Mbps) = 1,000,000 bits per second (bps)
- 1 gigabit per second (Gbps) = 1,000 (Mbps)

File sizes

- 1 megabyte (MB) = 1,024 kilobytes (KB)
- 1 gigabyte (GB) = 1,024 megabytes (MB) or 1,048,576 kilobytes (KB)

To change a traffic shaper's unit of measurement - CLI:

```
config firewall shaper traffic-shaper
    edit <traffic_shaper_name>
        set bandwidth-unit {kbps | mbps | gbps}
    end
```

Interface-based traffic shaping

You can enable traffic shaping on an interface. This allows you to enforce bandwidth limits for individual interfaces, by percentage. To configure interface-based traffic shaping, you must classify traffic in a traffic shaping policy, assign bandwidth percentages in a traffic shaping profile, and apply the traffic shaping profile as the egress traffic shaper on an interface.



Currently, only egress traffic shaping is available. To achieve ingress traffic shaping, you can typically configure egress traffic shaping on the alternate interface. For example, if you want to control inbound traffic on the WAN interface of the FortiGate, you can apply outbound traffic shaping to the LAN interface.

Classifying traffic

You can use a traffic shaping policy to classify traffic. Edit a traffic shaping policy using the `config firewall shaping-policy` command, and set the `class-id` command. A FortiGate stores the `class-id` on the kernel session, so that it can quickly categorize any traffic that matches the criteria you define in the traffic shaping policy.

To set the traffic class - CLI:

```
config firewall shaping-policy
  edit <shaping_policy_ID>
    ...
    set class-id <value>
  next
end
```

where `class-id` is a value in the range of 2 to 31.

Assigning bandwidth percentages

You can assign bandwidth percentages, using the `config firewall shaping-profile` command. Set a bandwidth guarantee using the `guaranteed-bandwidth-percentage` command and set a maximum bandwidth using the `maximum-bandwidth-percentage`.

To assign bandwidth percentages in a traffic shaping profile - CLI:

```
config firewall shaping-profile
  edit <egress_shaper_name>
    set default-class-id 2
    config shaping-entries
      edit 1
        set class-id 2
        set priority low {low | medium | high}
        set guaranteed-bandwidth-percentage 3
        set maximum-bandwidth-percentage 50
      next
      edit 3
        set class-id 5
        set priority low {low | medium | high}
        set guaranteed-bandwidth-percentage 3
        set maximum-bandwidth-percentage 50
      next
    end
  end
end
```

where you set the following variables:

Variable	Description
<code>default-class-id</code>	<p>The default class ID handles unclassified packets, including all local traffic. You must define the default class ID, since unclassified traffic must be controlled.</p> <p>Note that any traffic class that's defined in the traffic shaping policy, but isn't defined in the traffic shaping profile, is classified as part of the default class ID.</p>
<code>class-id</code>	The <code>class-id</code> is a value in the range of 2 to 31.
<code>priority</code>	The <code>priority</code> assigned (low, medium, high) to the class also plays a critical role in the bandwidth algorithm. Basically, priority decides which class can win when multiple classes compete for the available bandwidth on the interface.
<code>guaranteed-bandwidth-percentage</code>	<p>The <code>guaranteed-bandwidth-percentage</code> is a value in the range of 0 to 100 percent. The guaranteed bandwidth reserves a set amount of bandwidth for the class of traffic you select.</p> <p>For example, if you set the <code>guaranteed-bandwidth-percentage</code> to 3, then the FortiGate assigns at least 3% of the total bandwidth on the interface to that traffic class (as long as the current traffic volume of this class is more than 3% of the total volume). If the current traffic volume of this class is less than 3% of the total bandwidth of the interface, then it's not shaped.</p>
<code>maximum-bandwidth-percentage</code>	<p>The <code>maximum-bandwidth-percentage</code> is a value in the range of 0 to 100 percent. The maximum bandwidth defines the hard limit for traffic in that class. The class never has more bandwidth than the amount of bandwidth you define. You can assign 100% as the value, so that the class can potentially take all of the bandwidth of the designated interface.</p>



Important requirements:

- The `guaranteed-bandwidth-percentage` of the default class (in this example, class-id 2) must be greater than or equal to 1%. This ensures that local traffic always has some guaranteed bandwidth. However, the `guaranteed-bandwidth-percentage` of other classes can be 0.
- The `guaranteed-bandwidth-percentage` must not exceed the value of the `maximum-bandwidth-percentage`.
- The sum of `guaranteed-bandwidth-percentage` of all entries in one profile must not exceed 100%.

Apply the traffic shaping profile

You can apply the egress traffic shaper to an interface, using the `config system interface` command to edit the interface of your choice. Then, set the `inbandwidth` and `outbandwidth` values to the total amount of bandwidth that's available on the interface. Set the `egress-shaping-profile` to the traffic shaping profile you want to apply.

To apply the egress shaper to an interface - CLI:

```

config system interface
  edit <interface-name>
    set inbandwidth <limit>
    set outbandwidth <limit>
    set egress-shaping-profile <egress_shaper_name>
  next
end

```

where the `inbandwidth` and `outbandwidth` value is the total amount of bandwidth that's available on the interface, from a value in the range of 0 to 1677600 kbps.

You should set the `egress-shaping-profile` value to the traffic shaping profile you want to apply.

Example of competing priority classes

The following example can help you understand how the bandwidth algorithm uses both the class ID, and priority settings to determine which class wins when there are competing traffic classes. These examples are based on the assumption that the traffic volume of each class is larger than its allocated bandwidth.



NOTE: If a class has a small traffic volume, other classes can borrow unused bandwidth from it. In the following example, if class 2 has 100 MB of traffic and class 3 has 1 GB of traffic, then you should set the bandwidth for class 2 to 100 MB and for class 3 to 900 MB.

Class	Priority	guaranteed-bandwidth-percentage (%)	maximum-bandwidth-percentage (%)
2	high	20%	100%
3	low	20%	100%

If the profile configuration matches the table above, and the profile is applied to an egress interface with a total bandwidth of 1 GB, and both class 2 and class 3 have 1 GB of generated traffic, the results are the following:

Class	Priority	Actual bandwidth
2	high	80% of 1 GB (800 MB)
3	low	20% of 1 GB (200 MB)

The reason for the results are that both class 2 and 3 are assigned guaranteed bandwidth first, which is 200 MB each (20% of 1 GB). The remaining bandwidth of 600 MB is then allocated to class 2, because it has a higher priority.

The algorithm can get a bit more complex when you assign multiple classes with the same priority. When the same priority classes compete for available bandwidth, the allocation to each class is proportional to its `guaranteed-bandwidth-percentage`.

Here's a slightly more complex example:

Class	Priority	guaranteed-bandwidth-percentage (%)	maximum-bandwidth-percentage (%)
2	high	20%	100%
3	low	20%	100%
4	high	30%	100%

If the profile configuration matches the table above, and is attached to an egress interface with a total bandwidth of 1 GB, and classes 2, 3, and 4 have 1 GB of traffic generated, the results are the following:

Class	Priority	Actual bandwidth
2	high	200MB + 120MB = 320MB
3	low	200MB + 0 = 200MB
4	high	300MB + 180MB = 480MB

The reason for the results are that all classes are assigned the guaranteed bandwidth first, which is 200 MB, 200 MB, and 300 MB respectively. The remaining bandwidth of 300 MB is then allocated to class 2 and class 4, because of their higher priority settings. The allocation for the remaining 300MB is proportional to their guaranteed bandwidth. In this case, it is 120 MB for class 2 ($300 \text{ MB} * 20 / 50$) and 180MB for class 4 ($300 \text{ MB} * 30 / 50$).

Internet services support

The Internet Service Database (ISDB) and IP Reputation Database (IRDB) enhance traffic shaping criteria for traffic shaping policies.

To use Internet services in a traffic shaping policy, you must set the **Source** or **Destination** to one of the Internet services listed in the **Internet Service** tab. Internet service sources include **CloudServer-AWS**, **Proxy-Proxy.Server**, **SearchBot-Bing**, and more. Internet service destinations offer a wide range of services, such as **Github-Web**, **Salesforce-SMTP**, and **Netflix-DNS**.

The following image of the Traffic Shaping Policy page shows you where to find the **Internet Service** tab:

To create a traffic shaping policy that uses Internet services - GUI:

1. Create a new traffic shaping policy under **Policy & Objects > Traffic Shaping Policy**.
2. Add the **Internet Service** of your choice to the **Source** and **Destination** by selecting from the **Internet Service** tab on the far right.
3. Set the **Outgoing Interface** to the egress port that traffic passes through.

To create a traffic shaping policy that uses Internet services - CLI:

```
config firewall shaping-policy
  edit <shaping_policy_ID>
    set internet-service {enable | disable}
    set internet-service-id <service_ID>
    set internet-service-custom <custom_Internet_service_name>
    set internet-service-src {enable | disable}
    set internet-service-src-id <Internet_service_source_ID>
    set internet-service-src-custom <custom_Internet_service_source_name>
  next
end
```

where you set the following variables:

Option	Description
<code>internet-service</code>	Enables or disables the use of Internet services for this policy. If enabled, the FortiGate uses the Internet service destination address and service.

Option	Description
<code>internet-service-id</code>	The Internet service ID. For example: <ul style="list-style-type: none"> • 65536 Google-Others • 65537 Google-Web
<code>internet-service-custom</code>	Enter a custom Internet service name.
<code>internet-service-src</code>	Enables or disables the use of Internet services in source for this policy. If enabled, the FortiGate uses the Internet Services source address.
<code>internet-service-src-id</code>	The Internet service source ID. For example: <ul style="list-style-type: none"> • 65536 Google-Others • 65537 Google-Web
<code>internet-service-src-custom</code>	The custom Internet service source name. NOTE: This custom name must already be configured.

For more information about ISDB support in Firewall policies, see the ["Firewall policies"](#) on page 529 or *Networking Handbook*.

Differentiated services

Differentiated services (DiffServ) describes a set of end-to-end Quality of Service (QoS) capabilities. End-to-end QoS is the ability of a network to deliver service required by specific network traffic from one end of the network to another. By configuring differentiated services, you configure your network to deliver particular levels of service for different packets based on the QoS specified by each packet.

DiffServ is defined by RFC 2474 and 2475 as enhancements to IP networking to enable scalable service discrimination in the IP network without the need for per-flow state and signaling at every hop. Routers that can understand differentiated services sort IP traffic into classes by inspecting the DS field in IPv4 header or the Traffic Class field in the IPv6 header.

You can use the FortiGate DiffServ feature to change the DSCP (Differentiated Services Code Point) value for all packets accepted by a policy. The network can use these DSCP values to classify, mark, shape, and police traffic, and to perform intelligent queuing. DSCP features are applied to traffic by configuring the routers on your network to apply different service levels to packets depending on the DSCP value of the packet.

If the differentiated services feature isn't enabled, the FortiGate treats traffic as if the DSCP value is set to the default (00) and won't change the DSCP field of IP packets. DSCP values are also not applied to traffic if the traffic originates from a FortiGate itself.

The FortiGate applies the DSCP value and IPsec encryption to the differentiated services (formerly ToS) field in the first word of the IP header. The typical first word of an IP header, with the default DSCP value, is 4500:

- 4 for IPv4
- 5 for a length of five words
- 00 for the default DSCP value

You can change the packet's DSCP field for traffic initiating a session (forward) or for reply traffic (reverse) and enable each direction separately and configure it in the security policy.

Changes to DSCP values in a security policy effect new sessions. If traffic must use the new DSCP values immediately, clear all existing sessions.

To enable DSCP - CLI:


```
config firewall policy
  edit <policy_number>
    ...
    set diffserv-forward enable
    set diffservcode-forward <binary_integer>
    set diffserv-reverse enable
    set diffservcode-rev <binary_integer>
  end
```

For more information on the different DSCP commands, see the examples below and the CLI Reference. If you only set `diffserv-forward` and `diffserv-reverse` without setting the corresponding `diffservcode` values, the FortiGate resets the bits to zero.

For a list of DSCP values and their ToS equivalents see [Differentiated services on page 2811](#). DSCP values can also be defined within a shared traffic shaper as a single value, and per-IP traffic shaper for forward and reverse directions.

DSCP examples

For all the following DSCP examples, the FortiGate and client PC configuration is the following diagram and used firewall-based DSCP configurations.



Example

In this example, an ICMP ping is executed between User 1 and FortiGate B, through a FortiGate. DSCP is disabled on FortiGate B, and FortiGate A contains the following configuration:

```
config firewall policy
  edit 2
```

```

set srcintf port6
set dstintf port3
set src addr all
set dstaddr all
set action accept
set schedule always
set service ANY
set diffserv-forward enable
set diffservcode-forward 101110
end


```

As a result, FortiGate A changes the DSCP field for outgoing traffic, but not to its reply traffic. The binary DSCP values used map to the following hexadecimal.

ToS field values, which are observable by a sniffer (also known as a packet tracer):

- DSCP 000000 is TOS field 0x00
- DSCP 101110 is TOS field 0xb8, the recommended DSCP value for expedited forwarding (EF)

If you perform an ICMP ping between User 1 and User 2, the following output illustrates the IP headers for the request and the reply by sniffers on each of the network interfaces on the FortiGate. The right-most two digits of each IP header are the ToS field, which contains the DSCP value.



Example

In this example, an ICMP ping is executed between User 1 and FortiGate B, through FortiGate A. DSCP is disabled on FortiGate B, and FortiGate A contains the following configuration:

```


config firewall policy
edit 2
set srcintf port6
set dstintf port3
set src addr all
set dstaddr all
set action accept
set schedule always
set service ANY"
set diffserv-forward enable
set diffserv-rev enable
set diffservcode-forward 101110
set diffservcode-rev 101111
end

```

As a result, FortiGate A changes the DSCP field for both outgoing traffic and its reply traffic. The binary DSCP values in map to the following hexadecimal ToS field values, which are observable by a sniffer:

- DSCP 000000 is TOS field 0x**00**
- DSCP 101110 is TOS field 0x**b8**, the recommended DSCP value for expedited forwarding (EF)
- DSCP 101111 is TOS field 0x**bc**

If you perform an ICMP ping between User 1 and User 2, the output below illustrates the IP headers observed for the request and the reply by sniffers on each of the network interfaces on FortiGate A and FortiGate B. The right-most two digits of each IP header are the ToS field, which contains the DSCP value.



Example

In this example, an ICMP ping is executed between User 1 and FortiGate B, through FortiGate A. DSCP is enabled for both traffic directions on FortiGate A, and enabled only for reply traffic on FortiGate B. FortiGate A contains the following configuration:

```

config firewall policy
  edit 2
    set srcintf port6
    set dstintf port3
    set src addr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
    set diffserv-forward enable
    set diffserv-rev enable
    set diffservcode-forward 101110
    set diffservcode-rev 101111
  end

```

FortiGate B contains the following configuration:

```


config firewall policy
  edit 2
    set srcintf wan2
    set dstintf internal
    set src addr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
    set diffserv-rev enable
    set diffservcode-rev 101101
  end

```

As a result, FortiGate A changes the DSCP field for both outgoing traffic and its reply traffic, and FortiGate B changes the DSCP field only for reply traffic. The binary DSCP values in this configuration map to the following hexadecimal ToS field values:

- DSCP 000000 is TOS field **0x00**
- DSCP 101101 is TOS field **0xb4**
- DSCP 101110 is TOS field **0xb8**, the recommended DSCP value for expedited forwarding (EF)
- DSCP 101111 is TOS field **0xbc**

If you perform an ICMP ping between User 1 and User 2, the output below illustrates the IP headers observed for the request and the reply by sniffers on each of the network interfaces on FortiGate A and FortiGate B. The right-most two digits of each IP header are the ToS field, which contains the DSCP value.



Example

In this example, HTTPS and DNS traffic is sent from User 1 to FortiGate B, through FortiGate A. DSCP is enabled for both traffic directions on FortiGate A, and enabled only for reply traffic on FortiGate B. FortiGate A contains the following configuration:

```
config firewall policy
  edit 2
    set srcintf port6
    set dstintf port3
    set src addr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
    set diffserv-forward enable
    set diffserv-rev enable
    set diffservcode-forward 101110
    set diffservcode-rev 101111
  end
```


FortiGate B contains the following configuration:

```
config firewall policy
  edit 2
    set srcintf wan2
    set dstintf internal
    set src addr all
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
    set diffserv-rev enable
    set diffservcode-rev 101101
  end
```

As a result, FortiGate A changes the DSCP field for both outgoing traffic and its reply traffic, but FortiGate B changes the DSCP field only for reply traffic which passes through its internal interface. Since the example traffic doesn't pass through the internal interface, FortiGate B doesn't mark the packets. The binary DSCP values in this configuration map to the following hexadecimal ToS field values:

- DSCP 000000 is TOS field **0x00**
- DSCP 101101 is TOS field **0xb4**, which is configured on FortiGate B but not observed by the sniffer because the example traffic originates from the FortiGate itself, and therefore doesn't match that security policy.
- DSCP 101110 is TOS field **0xb8**, the recommended DSCP value for expedited forwarding (EF)
- DSCP 101111 is TOS field **0xbc**

If you sent HTTPS or DNS traffic from User 1 to FortiGate B, the following would illustrate the IP headers observed for the request and the reply by sniffers on each of FortiGate A's and FortiGate B's network interfaces. The right-most two digits of each IP header are the ToS field, which contains the DSCP value.



Traffic mapping

There are two types of traffic mapping: Type of Service (ToS) and Differentiated Services Code Point (DSCP). You can only use one method at a time. ToS is the default method. You can set the traffic mapping type using the CLI.

To configure traffic mapping - CLI:

```
config system global
    set traffic-priority {tos | dscp}
    set traffic-priority-level {low | medium | high }
end
```


Mapping of DSCP and ToS hexadecimal values for QoS


Service class	DSCP bits	DSCP value	ToS value	ToS hexadecimal
Network Control	111000	56-63	224	0xE0
Internetwork Control	110000	48-55	192	0xC0
Critical - Voice Data (RTP)	101110	46	184	0xB8
	101000	40	160	0xA0
Flash Override Video Data	100010	34	136	0x88
	100100	36	144	0x90
	100110	38	152	0x98
	100000	32	128	0x80
	011010	26	104	0x68
Flash Voice Control	011100	28	112	0x70
	011110	30	120	0x78
	011000	24	96	0x60


Service class	DSCP bits	DSCP value	ToS value	ToS hexadecimal
Immediate Deterministic (SNA)	010010	18	72	0x48
	010100	20	80	0x50
	010110	22	88	0x58
	010000	16	64	0x40
Priority Controlled Load	001010	10	40	0x28
	001100	12	48	0x30
	001110	14	56	0x38
	001000	8	32	0x20
Routine - Best Effort	000000	0	0	0x00
Routine - Penalty Box	000010	2	8	0x08


Traffic shaper monitor

You can view statistical information about traffic shapers and their bandwidth from FortiView > Traffic Shaping.


Refresh the information on the page.


Table View shows the following columns by default: Shaper, Bytes (Sent/Received), Sessions, Bandwidth, or Dropped Bytes. For more display options, right-click on the column header.


Bubble Chart shows you which resources consume the most bandwidth. Double-click on a traffic shaper to view more details. Determine whether more granular shaping is required by looking at the bandwidth usage by sources, destinations, applications, policies, and sessions.


FortiView Settings include the following options:

- Include Local traffic (Realtime Only)
- Include Unscanned Applications (Applications View Only)
- Auto update realtime visualizations
- Interval (seconds)
- Threat Weight Settings

Examples

While it is possible to configure QoS using a combination of security policies and ToS based priorities, and to distribute traffic over all six of the possible queues for each physical interface, the results of those configurations can be more difficult to analyze due to their complexity. In those cases, prioritization behavior can vary by several factors, including traffic volume, type of service (ToS) or differentiated services (DiffServ) markings, and correlation of session to a security policy.

The following simple examples illustrate QoS configurations using either prioritization by security policy, or prioritization by ToS bit, but not both. The examples also assume you are not configuring traffic shaping for interfaces that receive hardware acceleration from network processing units (NPU).

QoS using priority from security policies

Configurations implementing QoS using the priority values defined in the security policies are capable of applying bandwidth limits and guarantees.

In addition to configuring traffic shaping, you may also choose to limit the bandwidth accepted by each interface. This can be useful in scenarios where the bandwidth received on source interfaces frequently exceeds the maximum bandwidth limit defined in the security policy. Rather than waste processing power on packets that will get dropped later in the process, you may choose to preemptively police the traffic.

If you decide to implement QoS using security policies rather than ToS bit, the FortiGate applies QoS to all packets controlled by the policy. This type of control is less granular than prioritization by ToS bit, but has the benefits of correlating quality of service to a security policy. This correlation enables you to distribute traffic over up to four of the possible 6 priority queues (queue 0 to queue 3), doesn't require other devices in your network to set or respect the ToS bit, and enables you to configure bandwidth limits and guarantees.

In the following example, we limit the bandwidth accepted by each source interface, limit the bandwidth used by sessions controlled by the security policy, and then configure prioritized queuing on the destination interface based upon the priority in the security policy, subject to alternative assignment to queue 0 when necessary to achieve the guaranteed packet rate.

To limit bandwidth accepted by an interface - CLI:

```
config system interface
  edit <interface_name>
    set inbandwidth <limit>
  next
end
```

where <rate_int> is the bandwidth limit in Kbps. Excess packets are dropped.

To configure bandwidth guarantees, limits, and priorities - GUI:

1. Go to **Policy & Objects > Traffic Shapers** and select the Create New + sign.
2. Select **Shared** or **Per-IP**.
3. Enter a name for the traffic shaper.

4. Select the **Traffic Priority**.
High has a priority value of 1, Medium is 2, and Low is 3. While the current packet rate is below Guaranteed Bandwidth, the FortiGate disregards this setting, and instead uses priority queue.
5. Enable **Max Bandwidth** and enter a value.
Packets greater than this rate are discarded.
6. Enable **Guaranteed Bandwidth** and enter a value, if any.
Bandwidth guarantees affect prioritization. While packet rates are less than this rate, they use priority queue 0. If this isn't the effect you intend, consider entering a small guaranteed rate, or enter 0 to effectively disable bandwidth guarantees.
7. Enable **DSCP** and set a value.
8. Select **OK**.



Per-IP shapers also include the option to set a maximum number of concurrent connections and to set both **Forward DSCP** and **Reverse DSCP**.

Sample configuration

This sample configuration limits ingressing bandwidth to 500 Kbps. It also applies separate traffic shapers to FTP and HTTP traffic. In addition to the interface bandwidth limit, HTTP traffic is subject to a security policy bandwidth limit of 200 Kbps.

All egressing FTP traffic greater than 10 Kbps is subject to a low priority queue (queue 3), while all egressing HTTP traffic greater than 100 Kbps is subject to a medium priority queue (queue 2). That is, unless FTP traffic rates are lower than their guaranteed rate, and web traffic rates are greater than their guaranteed rate, FTP traffic is lower priority than web traffic.

Traffic less than these guaranteed bandwidth rates use the highest priority queue (queue 0).

To set the inbandwidth limits - CLI:

This setting is only available in the CLI.

```
config system interface
  edit wan1
    set inbandwidth 500
  next
end
```

Create traffic shapers for FTP and HTTP.

To configure an FTP shaper - GUI:

1. Go to **Policy & Objects > Traffic Shapers**, and select the **Create New** "Plus" icon.
2. Select **Shared**.
3. Enter **FTP** for the name of the traffic shaper.
4. Set **Traffic Priority** to **Low**.
5. Select the **Guaranteed Bandwidth** checkbox and enter **10** Kbps.
6. Select the **Maximum Bandwidth** checkbox and enter **500** Kbps.

7. Select **OK**.
8. Select the FTP traffic shaper, right-click it, and select **Edit in CLI**. Type the following command:

```
set per-policy
end
```

To configure an HTTP shaper - GUI:

1. Select the **Create New** "Plus" icon.
2. Set **Type** to **Shared**.
3. Enter **HTTP** for the name of the traffic shaper.
4. Set **Traffic Priority** to **Medium**.
5. Select the **Guaranteed Bandwidth** checkbox and enter **100 Kbps**.
6. Select the **Maximum Bandwidth** checkbox and enter **200 Kbps**.
7. Select **OK**.
8. Select the HTTP traffic shaper, right-click it, and select **Edit in CLI**. Type the following command:

```
set per-policy
end
```

To add the FTP shaper to a traffic shaping policy - GUI:

1. Go to **Policy & Objects > Traffic Shaping Policy** and click **Create New** to create a traffic shaping policy for FTP.
2. Set the **Matching Criteria** to the following:

Source	all
Destination address	all
Service	FTP

3. Under **Apply shaper**, set the following:

Outgoing interface	any (The outgoing interface should match the outgoing interface of the security policy you wish to apply shaping to.)
Shared Shaper	Enable Shared Shaper and select FTP from the drop-down menu.
Reverse Shaper	Enable Shared Shaper and select FTP from the drop-down menu.
Enable this policy	Enable this policy.

4. Select **OK**.

To add the HTTP shaper to a traffic shaping policy - GUI:

1. Go to **Policy & Objects > Traffic Shaping Policy** and click **Create New** to create a traffic shaping policy for HTTP.
2. Set the **Matching Criteria** to the following:

Source	all
Destination address	all
Service	HTTP

3. Under **Apply shaper**, set the following:

Outgoing interface	any (The outgoing interface should match the outgoing interface of the security policy you want to apply traffic shaping to.)
Shared Shaper	Enable Shared Shaper and select HTTP from the drop-down menu.
Reverse Shaper	Enable Shared Shaper and select HTTP from the drop-down menu.
Enable this policy	Enable this policy.

4. Select **OK**.
5. On the policy list page, move the FTP traffic shaping policy to the top of the list by clicking the far left column to drag and drop it. The HTTP traffic shaping policy should be below the FTP policy, and more general Internet access policies should be at the bottom of the policy list.

To configure the FTP and HTTP traffic shapers - CLI:

```
config firewall shaper traffic-shaper
  edit FTP
    set maximum-bandwidth 500
    set guaranteed-bandwidth 10
    set per-policy enable
    set priority low
  next
  edit HTTP
    set maximum-bandwidth 200
    set guaranteed-bandwidth 100
    set per-policy enable
    set priority medium
  end
```

To add each traffic shaper to a traffic shaping policy - CLI:

```
config firewall shaping-policy
  edit 1 <shaping_policy_ID_number>
    set srcaddr all
    set dstaddr all
    set service ALL
    set dstintf wan1 <outgoing_interface>
    set traffic-shaper FTP
  next
  edit 2 <shaping_policy_ID_number>
    set srcaddr all
    set dstaddr all
    set service ALL
    set dstintf wan1 <outgoing_interface>
    set traffic-shaper HTTP
```

```
next
move 1 before 2
end
```

QoS using priority from ToS or DiffServ

Configurations implementing QoS using the priority values defined in either global or specific ToS bit values are not capable of applying bandwidth limits and guarantees, but are capable of prioritizing traffic at per-packet levels, rather than uniformly to all services matched by the security policy.

In addition to configuring traffic prioritization, you may also choose to limit bandwidth that's received by each interface. This can sometimes be useful in scenarios where you want to limit traffic levels, but don't want to configure traffic shaping within a security policy. This has the benefit of policing traffic at a point before the FortiGate performs most processing.

Note that if you implement QoS using ToS octet rather than security policies, the FortiGate applies QoS on a packet-by-packet basis, and priorities may be different for packets and services controlled by the same security policy. This is more granular control than prioritization by security policies, but has the drawbacks that quality of service may not be uniform for multiple services controlled by the same security policy, packets only use up to three of the six possible queues (queue 0 to queue 2), and bandwidth can't be guaranteed. Other devices in your network must also be able to set or preserve ToS bits.

In this example, we limit the bandwidth accepted by each source interface, and then configure prioritized queuing on the destination interface based upon the value of the ToS bit located in the IP header of each accepted packet.

To limit bandwidth accepted by an interface - CLI:

```
config system interface
  edit <interface_name>
    set inbandwidth <limit>
  next
end
```

where <limit> is the bandwidth limit in Kbps. Excess packets are dropped.

To configure the global priority value - CLI:

```
config system global
  set tos-based-priority {high | low | medium}
end
```

where `high` has a priority value of 0 and `low` is 2.

If you want to prioritize some ToS bit values differently than the global ToS-based priority, configure the priority for packets with that ToS bit value using the following commands:

```
config system tos-based-priority
  edit <id_int>
    set tos [0-15]
    set priority {high | low | medium}
  next
end
```

where `tos` is the value of the ToS bit in the packet's IP header, and `high` has a priority value of 0 and `low` is 2. Priority values configured in this location will override the global ToS-based priority.

Sample configuration

This sample configuration limits ingress bandwidth to 500 Kbps. It also queues egress traffic based upon the ToS bit in the IP header of ingress packets.

Unless specified for the packet's ToS bit value, packets use the low priority queue (queue 2). For ToS bit values 4 and 15, the priorities are specified as medium (value 1) and high (value 0), respectively.

```
config system interface
  edit wan1
    set inbandwidth 500
  next
end
config system global
  set tos-based-priority low
end
config system tos-based-priority
  edit 4
    set tos 4
    set priority medium
  next
  edit 15
    set tos 15
    set priority high
  next
end
```

Example setup for VoIP

In this example, there are three traffic shaping requirements for a network:

- Voice over IP (VoIP) requires a guaranteed, high-priority for bandwidth for telephone communications.
- FTP bursts must be contained so it doesn't consume any available bandwidth. As such, this traffic needs to be throttled to a smaller amount.
- A consistent bandwidth requirement is needed for all other email and web-based traffic.

To enable this requirement, you need to create three separate traffic shapers and three traffic shaping policies for each traffic type.

In this example, the values used aren't recommended values.

Creating the traffic shapers

First create the traffic shapers that define the maximum and guaranteed bandwidth. The shared traffic shapers are used with some applied per-policy and some applied to all policies, to better control traffic.

VoIP shaper

The VoIP functionality is a key component to the business as a communication tool and as such requires a guaranteed bandwidth. This traffic shaper is a high priority traffic shaper.

To create a VoIP shaper - GUI:

1. Go to **Policy & Objects > Traffic Shapers** and select **Create New**.
2. Set the **Type** to **Shared**.
3. Enter the **Name** `voip`.
4. Set the **Traffic Priority** to **High**.
5. Select **Maximum Bandwidth** and enter `1000 Kbps`.
6. Select **Guaranteed Bandwidth** and enter `800 Kbps`.
7. Select **OK**.
8. Select the HTTP shaper, right-click it, and select **Edit in CLI**. Type the following command:

```
set per-policy
end
```

To create a VoIP shaper - CLI:

```
config firewall shaper traffic-shaper
edit voip
    set maximum-bandwidth 1000
    set guaranteed-bandwidth 800
    set per-policy enable
    set priority high
end
```

Setting the traffic shaper to **per-policy** ensures that regardless of the number of policies that use this traffic shaper, the defined bandwidth is always the same. At the same time, the bandwidth is continually guaranteed at 800 Kbps but, if available, can be as much as 1000 Kbps. Setting the priority to high ensures that the FortiGate considers VoIP traffic the most important.

FTP traffic shaper

The FTP traffic shaper sets the maximum bandwidth to use to avoid sudden spikes by sudden uploading or downloading of large files, and interfering with other more important traffic.

To create a FTP shaper - GUI:

1. Go to **Policy & Objects > Traffic Shapers** and **Create New**.
2. Set the **Type** to **Shared**.
3. Enter the **Name** `ftp`.
4. Set the **Traffic Priority** to **Low**.
5. Select **Maximum Bandwidth** and enter `200 Kbps`.
6. Select **Guaranteed Bandwidth** and enter `200 Kbps`.
7. Select **OK**.

To create a FTP shaper - CLI:

```
config firewall shaper traffic-shaper
edit ftp
    set maximum-bandwidth 200
```

```
set guaranteed-bandwidth 200
set priority low
end
```

For this traffic shaper, the maximum and guaranteed bandwidth are set to a low value and to the same value. In this case, the bandwidth is restricted to a specific amount. Setting the traffic priority to a low value ensures that more important traffic passes before FTP traffic.

Regular traffic shaper

The regular traffic shaper sets the maximum bandwidth and guaranteed bandwidth for everyday business traffic such as web and email traffic.

To create a regular traffic shaper - GUI:

1. Go to **Policy & Objects > Traffic Shapers** and **Create New**.
2. Set the **Type** to **Shared**.
3. Enter the **Name** `daily_traffic..`
4. Set the **Traffic Priority** to **Medium**.
5. Select **Maximum Bandwidth** and enter `600 Kbps`
6. Select **Guaranteed Bandwidth** and enter `600 Kbps`.
7. Select **OK**.

To create a regular traffic shaper - CLI:

```
config firewall shaper traffic-shaper
edit daily_traffic
set maximum-bandwidth 600
set guaranteed-bandwidth 600
set per-policy enable
set priority medium
end
```

For this traffic shaper, the maximum and guaranteed bandwidth are set to a moderate value of 600 Kbps. It's also set per policy, which ensures each security policy for day-to-day business traffic has the same distribution of bandwidth.

Creating traffic shaping policies

To employ the traffic shapers, create traffic shaping policies that apply to your existing security policy. Create a separate policy for each service and apply the traffic shaper to the outgoing interface you want to use. For example, a policy for FTP traffic, a policy for SIP, and so on.

For the following steps, the VoIP traffic shaper is enabled as well as the reverse direction. This ensures that return traffic for a VoIP call has the same guaranteed bandwidth as the outgoing call. The example below shows how to enable each traffic shaper in a traffic shaping policy.

In this example, the traffic shaping policies will apply traffic shaping to the following security policy:

Incoming interface	lan (Internal interface)
---------------------------	--------------------------

Source address	All
Outgoing interface	WAN1
Destination address	All
Schedule	always
Service	all
Action	ACCEPT

To create a VOIP traffic shaping policy- GUI:

1. Go to **Policy & Objects > Traffic Shaping Policy** and select **Create New**.
2. Now create a traffic shaping policy that matches the settings you entered for the security policy:

Source	All
Destination	All
Service	All
Application Category	VoIP
Application	SIP
URL Category	Internet Telephony
Outgoing Interface	wan1

3. Enable **Shared Shaper**, select the VoIP traffic shaper that you created in the previous steps.
4. Enable **Reverse Shaper**, select the VoIP traffic shaper that you created in the previous steps.
5. Select **Enable this policy**.
6. Select **OK**.

To create a VOIP traffic shaping policy - CLI:

```
config firewall shaping-policy
edit 1 <shaping_policy_ID_number>
set srcaddr all
set dstaddr all
set service ALL
set application 34640 <SIP>
set app-category 3 <VoIP>
set url-category 76 <Internet Telephony>
set dstintf wan1 <outgoing_interface>
set traffic-shaper voip <high_priority_custom_shaper>
set reverse-traffic-shaper voip <high_priority_custom_shaper>
end
```

To create an FTP traffic shaping policy - GUI:

1. Go to **Policy & Objects > Traffic Shaping Policy** and select **Create New**.
2. Now create a traffic shaping policy that matches the settings you entered for your security policy:

Source	All
Destination	All
Service	FTP
Outgoing Interface	wan1

3. Enable **Shared Shaper**, select the FTP shaper created in the previous steps.
4. Enable **Reverse Shaper**, select the FTP shaper created in the previous steps.
5. Select **Enable this policy**.
6. Select **OK**.

To create an FTP traffic shaping policy - CLI:

```
config firewall shaping-policy
edit 2 <shaping_policy_ID_number>
set srcaddr all
set dstaddr all
set service FTP
set dstintf wan1 <outgoing_interface>
set traffic-shaper FTP <low_priority_custom_shaper>
set reverse-traffic-shaper FTP <low_priority_custom_shaper>
end
```

To create a regular traffic shaping policy - GUI:

1. Go to **Policy & Objects > Traffic Shaping Policy** and select **Create New**.
2. Now create a traffic shaping policy that matches the settings you entered for your security policy:

Source	All
Destination	All
Service	ALL
Outgoing Interface	wan1

3. Enable **Shared Shaper**, select the medium-priority shaper.
4. Enable **Reverse Shaper**, select the medium-priority shaper.
5. Select **Enable this policy**.
6. Select **OK**.

To create a regular traffic shaping policy - CLI:

```
config firewall shaping-policy
```

```
edit 3 <shaping_policy_ID_number>
  set srcaddr all
  set dstaddr all
  set service ALL
  set dstintf wan1 <outgoing_interface>
  set traffic-shaper medium-priority <default_shaper>
  set reverse-traffic-shaper medium-priority <default_shaper>
end
```

To order your traffic shaping policies - CLI:

```
config firewall shaping-policy
  move 1 before 2
  move 3 below 2
end
```



Ensure that your high priority SIP/VoIP policy is at the top of the policy list, the low priority FTP traffic shaper comes second, and the medium priority regular traffic shaper comes last. Restrictive policies should always go above more general access policies.

Alternate method of enabling traffic shaping in the security policy

It's also possible to create three separate security policies for each type of traffic (VoIP, FTP, and regular). You can enable traffic shaping individually within each security policy in the CLI only, like the example shown below:

To enable traffic shaping in the security policy - CLI:

```
config firewall policy
  edit 6
    set srcintf <internal_interface>
    set scraddr all
    set dstintf wan1
    set dstaddr all
    set action accept
    set schedule always
    set service sip
    set traffic-shaper voip
    set reverse-traffic-shaper voip
  end
```

Troubleshooting traffic shaping

This chapter outlines some troubleshooting tips and steps to diagnose the traffic shapers and whether they're working correctly. These diagnose commands include:

- `diagnose system tos-based-priority`
- `diagnose firewall shaper traffic-shaper`
- `diagnose firewall per-ip-shaper`
- `diagnose debug flow`

Interface diagnosis

To optimize traffic shaping performance, first ensure that the network interface's Ethernet statistics are clean of errors, collisions, or buffer overruns. To check the interface, enter the following diagnose command to see the traffic statistics:

```
diagnose hardware deviceinfo nic <port_name>
```

Traffic shaper diagnose commands

There are specific diagnose commands you can use to verify the configuration and flow of traffic, including packet loss due to the employed traffic shaper.

All of these diagnose troubleshooting commands are supported in both IPv4 and IPv6.

ToS command

Use the following command to list command to view information of the ToS lists and traffic:

```
diagnose system tos-based-priority
```

This example displays the priority value currently correlated with each possible ToS bit value. Priority values are displayed in order of their corresponding ToS bit values, which can range between 0 and 15, from lowest ToS bit value to highest.

For example, if you configured ToS-based priorities, the following appears:

```
0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0
```

This reflects that all packets are currently using the same default priority, high (value 0).

If you configured a ToS-based priority of `low` (value 2) for packets with a ToS bit value of 3, the following appears:

```
0 0 0 2 0 0 0 0 0 0 0 0 0 0 0 0
```

This reflects that most packets are using the default priority value, except those with a ToS bit value of 3.

Shared traffic shaper

To view information for the shared traffic shaper for security policies, enter the command:

```
diagnose firewall shaper traffic-shaper list
```

The resulting output displays the information on all available traffic shapers. The more traffic shapers that are available, the longer the list. For example:

```
name Throughput
maximum-bandwidth 1200000 Kb/sec
guaranteed-bandwidth 50000 Kb/sec
current-bandwidth 0 B/sec
priority 1
packets dropped 0
```

Additional commands include:

`diagnose firewall shaper traffic-shaper state` - provides the total number of traffic shapers on the FortiGate.

`diagnose firewall shaper traffic-shaper stats` - provides summary statistics on the shapers.

Sample output looks like the following:

```
shapers 9 ipv4 0 ipv6 0 drops 0
```

Per-IP traffic shaper

To view information for the per-IP traffic shaper for security policies, enter the command:

```
diagnose firewall shaper per-ip-shaper list
```

The resulting output displays the information on all available per-IP traffic shapers. The more traffic shapers that are available, the longer the list. For example:

```
name accounting_group
maximum-bandwidth 200000 Kb/sec
maximum-concurrent-session 55
packet dropped 0
```

Additional commands include:

`diagnose firewall shaper per-ip-shaper state` - provides the total number of per-ip traffic shapers on the FortiGate.

`diagnose firewall shaper per-ip-shaper stats` - provides summary statistics on the traffic shapers. Sample output looks like the following:

```
memory allocated 3 packet dropped: 0
```

You can also clear the per-ip statistical data to begin a fresh diagnosis using:

```
diagnose firewall shaper per-ip-shaper clear
```

Packet loss with statistics on traffic shapers

For each traffic shaper there are counters that allow you to verify if packets have been discarded. To view this information, enter the `diagnose firewall shaper` command in the CLI. The results look similar to the following output:

```
diagnose firewall shaper traffic-shaper list
name limit_GB_25_MB_50_LQ
maximum-bandwidth 50 Kb/sec
guaranteed-bandwidth 25 Kb/sec
current-bandwidth 51 Kb/sec
priority 3
dropped 1291985
```

The diagnose command output is different if the diagnose firewall shapershapers are configured either per-policy or shared between policies.

For per-IP the output is:

```
diagnose firewall shaper per-ip-shaper list

name accounting_group
maximum-bandwidth 200000 Kb/sec
maximum-concurrent-session 55
packet dropped 3264220
```

Packet lost with the debug flow

When using the debug flow diagnostic command, there is a specific message information that a packet has exceed the diagnose firewall shapershaper limits and therefore discarded:

```
diagnose debug flow show console enable
diagnose debug flow filter addr 10.143.0.5
diagnose debug flow trace start 1000

id=20085 trace_id=11 msg="vd-root received a packet(proto=17, 10.141.0.11:3735-
>10.143.0.5:5001) from port5."
id=20085 trace_id=11 msg="Find an existing session, id=0000eabc, original direction"
id=20085 trace_id=11 msg="exceeded shaper limit, drop"
```

Session list details with dual traffic shaper

When a security policy has a different traffic shaper for each direction, it's reflected in the session list output from the CLI:

```
diagnose system session list

session info: proto=6 proto_state=02 expire=115 timeout=3600 flags=00000000 sock
flag=00000000 sockport=0 av_idx=0 use=4
origin-shaper=Limit_25Mbps prio=1 guarantee 25600/sec max 204800/sec traffic 48/sec
reply-shaper=Limit_100Mbps prio=1 guarantee 102400/sec max 204800/sec traffic 0/sec
ha_id=0 hakey=44020
policy_dir=0 tunnel=/
state=may_dirty rem os rs
statistic(bits/packets/allow_err): org=96/2/1 reply=0/0/0 tuples=2
orgin->sink: org pre->post, reply pre->post dev=2->3/3->2 gwy=10.160.0.1/0.0.0.0
hook=pre dir=org act=dnat 192.168.171.243:2538->192.168.182.110:80(10.160.0.1:80)
hook=post dir=reply act=snat 10.160.0.1:80->192.168.171.243:2538(192.168.182.110:80)
pos/(before,after) 0/(0,0), 0/(0,0)
misc=0 policy_id=2 auth_info=0 chk_client_info=0 vd=0 serial=00011e81 tos=ff/ff app=0
dd_type=0 dd_rule_id=0
```

Additional information

- Packets discarded by the traffic shaper impact flow-control mechanisms like TCP. For more accurate testing results, use the UDP protocol.
- Traffic shaping accuracy is optimum for security policies without a protection profile where no FortiGate content inspection is processed.

- Don't oversubscribe an outbandwidth throughput. For example, $\text{sum}[\text{guaranteed BW}] < \text{outbandwidth}$. For accuracy in bandwidth calculation, it's required to set the "outbandwidth" parameter on the interfaces. For more information, see ["Bandwidth guarantee, limit, and priority interactions" on page 2784](#).
- The FortiGate isn't prioritizing traffic based on the DSCP marking configured in the security policy. However, ToS-based prioritizing can be made at ingress. For more information, see ["Traffic shaping methods" on page 2791](#).

Chapter 26 - Transparent Mode

- [Transparent mode overview](#): an overview of transparent mode, including available features.
- [Installation](#): instructions for installing a FortiGate in transparent mode.
- [Networking in transparent mode](#): how networking is configured in transparent mode.
- [Firewalls and security in transparent mode](#): information about using firewalls and security scanning in transparent mode.
- [IPsec VPN in transparent mode](#): configuring IPsec VPNs using FortiGates in transparent mode.
- [Using FortiManager and FortiAnalyzer](#): using external management and logging.
- [High availability in transparent mode](#): configuring transparent FortiGates in HA mode.
- [Best practices](#): the best practices for using transparent mode.

What's new in FortiOS 6.0

The following list contains new transparent mode features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- ["Source prefixes" on page 2841](#)

Transparent mode overview

This section contains an overview of transparent mode. It contains the following topics:

- [What is transparent mode?](#)
- [Transparent mode features](#)

What is transparent mode?

A FortiGate can operate in one of two modes: transparent or NAT/route mode.

In transparent mode, the FortiGate is installed between the internal network and the router. In this mode, the FortiGate does not make any changes to IP addresses and only applies security scanning to traffic. When a FortiGate is added to a network in transparent mode, no network changes are required, except to provide the FortiGate with a management IP address. transparent mode is used primarily when there is a need to increase network protection but changing the configuration of the network itself is impractical.

In NAT/route mode, a FortiGate is installed as a gateway or router between two networks. This allows the FortiGate to hide the IP addresses of the private network using network address translation (NAT).

Transparent mode features

Different FortiOS features are available depending on whether your FortiGate is in transparent or NAT/route mode. The following table shows which features are available for each mode.



For a FortiGate in transparent mode, the maximum number of Interfaces per VDOM is 254. This value includes both physical and virtual interfaces.

For any other maximum values, please consult the Maximum Values Table, available at docs.fortinet.com.

Feature	NAT	Transparent	Comment
Unicast routing/policy-based routing	Yes	No	
VIP/IP pools/ NAT	Yes	Yes	Configurable from CLI only in transparent mode
Multicast routing	Yes	No	Options are available to forward multicast packets
L2 forwarding	No	Yes	In transparent mode, other frames than IP can be forwarded, but only without security scanning

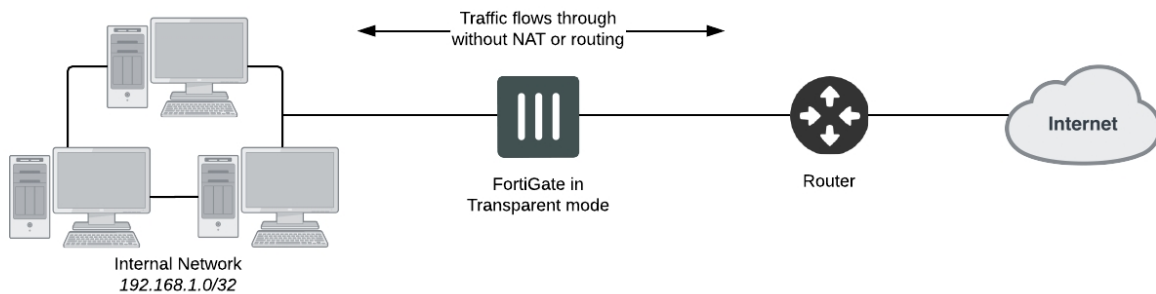
Feature	NAT	Transparent	Comment
Firewall (packet filtering/NAT/Authentication)	Yes	Yes	
IPv6 capable	Yes	Yes	
Traffic shaping (type of service)	Yes	Yes	
Hardware acceleration	Yes	Yes	
All security profile features (ex IPS, Application Control, Web Filtering, etc ...)	Yes	Yes	
Security Fabric	Yes	No	
FortiView	Yes	Yes	
IPsec VPN	Yes	Yes	Only policy based IPsec VPNs are supported for transparent mode
SSL VPN	Yes	No	
High availability (HA) - virtual cluster	Yes	Yes	
802.3ad (LACP/port aggregation)	Yes	Yes	
HA port redundancy	Yes	Yes	FortiGate hardware dependent
802.1q - VLAN trunking	Yes	Yes	
802.1d - spanning tree	No	No	Option to forward VPDUs
Logging and reporting (disk and memory logging, FortiCloud, syslog, and FortiAnalyzer)	Yes	Yes	
Managed by FortiManager	Yes	Yes	

Installation

This section contains information about installing a FortiGate in transparent mode. It contains the following topics:

- [Installing a FortiGate in transparent mode](#)
- [Using a virtual wire pair to simplify transparent mode](#)
- [Management IP configuration](#)

Installing a FortiGate in transparent mode



Changing to transparent mode removes most configuration changes made in NAT/route mode. To keep your current NAT/route mode configuration, backup the configuration using the **System Information** widget, found in the **Dashboard**.

1. Before connecting the FortiGate to your network, go to the **Dashboard** and locate enter the following command into the **CLI Console**:

```
config system settings
  set opmode transparent
  set manageip <address and netmask>
  set gateway <address>
end
```

2. Access the web-based manager by browsing to the new management IP.
3. (Optional) The FortiGate's DNS Settings are set to use FortiGuard DNS servers by default, which is sufficient for most networks. However, if you need to change the DNS servers, go to **Network > DNS** and add **Primary** and **Secondary** DNS servers. Select **Apply**.
4. If your network uses IPv4 addresses, go to **Policy & Objects > IPv4 Policy** and select **Create New** to add a security policy that allows users on the private network to access the Internet.

If your network uses IPv6 addresses, go to **Policy & Objects > IPv6 Policy** and select **Create New** to add a

security policy that allows users on the private network to access the Internet. If the IPv6 menu option is not available, go to **System > Feature Visibility**, turn on **IPv6**, and select **Apply**. For more information on IPv6 networks, see the IPv6 Handbook.

Set the **Incoming Interface** to the internal interface and the **Outgoing Interface** to the Internet-facing interface (typically WAN1). You will also need to set **Source Address**, **Destination Address**, **Schedule**, and **Service** according to your network requirements. You can set these fields to the default all/ANY settings for now but should create the appropriate objects later after the policies have been verified.

5. Make sure the **Action** is set to **ACCEPT**. Select **OK**.



It is recommended to avoid using any security profiles, such as AntiVirus or web filtering, until after you have successfully installed the FortiGate. After the installation is verified, you can apply any required security profiles.

For more information about using security profiles, see the Security Profiles handbook.

6. Go to the **Dashboard** and locate the **System Resources** widget. Select **Shutdown** to power off the FortiGate. Alternatively, you can also use the CLI command `execute shutdown`.
7. Connect the FortiGate between the internal network and the router.
8. Connect the Internet-facing interface to the router's internal interface and connect the internal network to the FortiGate using an internal port (typically port 1).
9. Power on the FortiGate. You will experience downtime before the FortiGate starts up completely.


Results

Users on the internal network are now able to browse to the Internet. They should also be able to connect to the Internet using any other protocol or connection method that you defined in the security policy.



If a FortiGate operating in transparent mode is installed between your internet network and a server that is providing a network service to the internal network, such as DNS or DHCP, you must add a wan1-to-internal policy to allow the server's response to flow through the FortiGate and reach the internal network.

Using a virtual wire pair to simplify transparent mode



A virtual wire pair consists of two interfaces that have no IP addresses and all traffic received by one interface in the pair can only be forwarded out the other; as controlled by firewall policies. Since the interfaces do not have IP addresses, you can insert a virtual wire pair into a network without having to make any changes to the network.



Interfaces used in a virtual wire pair cannot be used for admin access to the ISFW FortiGate. Before creating a virtual wire pair, make sure you have a different port configured to allow admin access using your preferred protocol.

1. Go to **Network > Interfaces** and select **Create New > Virtual Wire Pair**. Add two ports to the virtual wire pair. These ports cannot be part of a switch, such as the default **internal/lan** interface.
2. Go to **Policy & Objects > IPv4 Virtual Wire Pair Policy** and create a policy will allow traffic to flow between the two ports. Give the policy an appropriate **Name**. Select the direction that traffic is allowed to flow. Configure the other firewall options as needed.
3. If necessary, create a second virtual wire pair policy allowing traffic to flow between the ports in the opposite direction.

Traffic can now flow between the two ports. Go to **FortiView > All Segments > Policies** to see traffic flowing through both policies.

Management IP configuration

A FortiGate in transparent mode can be assigned with a single IP address for remote access management and multiple static routes can be configured. This can be used if in-band management wants to be applied.

When out-of-band management is desired (dedicated interface for remote management access), it is recommended to use a separate VDOM in NAT/route mode.

In-band management details and example

The management IP address is bound to all ports or VLANs belonging to the same VDOM. Remote access services are subject to the same rules as in NAT/route mode, and must be enabled/disabled on each port.

Example of management IP configuration in transparent mode:

```
config system settings
  set manageip 10.1.1.100/255.255.255.0
end
config router static
  edit 1
    set gateway 10.1.1.254
  next
end
config system interface
  edit port1
    set allowaccess ping ssh https snmp
  end
```

It is also possible to add a second IP address for management and additional default routes:

```
config system settings
  set opmode transparent
  set manageip 192.168.182.136/255.255.254.0 10.1.1.1/255.255.255.0
```

```
end

config router static
  edit 1
    set gateway 192.168.183.254
  next
  edit 2
    set gateway 10.1.1.254
  next
end
```



`ping-server` (dead gateway detection) is not supported in transparent mode.

Out-of-band management details and example

When VDOM is enabled and the VDOMs are operating in transparent mode, it is recommended, to avoid L2 loops and allow more routing flexibility, to keep one VDOM (generally the root VDOM) in NAT/route mode, with one or more VLAN or physical interface as out-of-band management.



The management VDOM must have IP connectivity to the Internet to allow communication with the FDS and retrieve services information (antivirus, IPS, FortiGuard, FortiCare, etc...). All syslog and FortiManager communication also go through the management VDOM.

Networking in transparent mode

This section contains information about networking concepts in transparent mode. It contains the following topics:

- [Static routing](#)
- [Packet forwarding](#)
- [Network address translation \(NAT\)](#)
- [VLANs and forwarding domains](#)
- [Inter-VDOM links between NAT/route and transparent VDOMs](#)
- [Packet forwarding using Cisco protocols](#)
- [Configuration example](#)

Static routing

The following sections include information about configuring static routing in transparent mode:

- [Overview](#)
- [Source prefixes](#)

Overview

When you configure routing in transparent mode on a FortiGate, all interfaces must be connected to the same subnet. This means all traffic comes from and leaves on the same subnet. This is important because it limits the static routing options to only gateways that are attached to this subnet. For example, if you have only one router that connects your network to the Internet, all static routing on the FortiGate uses this gateway. For this reason, static routing on a FortiGate in transparent mode may be a bit different, but it's not as complex as routing in NAT mode.

To view the routing table in transparent mode, go to **Network > Static Routes**. When you view entries for static routes in transparent mode, you'll see the following settings:

Field	Description
Destination	When Subnet is selected, shows the IP address and netmask of the destination of the traffic being routed. 0.0.0.0 is the default route and matches all traffic destinations.
Gateway	Specifies the IP address of the next hop for traffic. This is usually the IP address of a router on the edge of your network.

Field	Description
Priority	<p>The FortiGate uses the priority if there's more than one match for a route. This allows you to use multiple routes, but configure preferred routes.</p> <p>Routes with a larger value have a lower priority. If the preferred route isn't available, another route is used instead. If there is more than one match for a route, and the routes have the same priority, the FortiGate uses Equal Cost Multiple Path (ECMP) to share traffic between the routes.</p> <p>The possible values are 0 to 4294967295. This setting only applies to static routes. The priority for routes that are dynamically learned from routing protocols is 0.</p>

Source prefixes

If a FortiGate has more than one management IP address and default route, packets can't differentiate between them and may reach the wrong management IP address. To avoid this, you can configure a source prefix that allows the FortiGate to differentiate between multiple default routes.

To configure source prefixes - CLI

```
config router static
  edit <sequence number>
    set gateway <IP address>
    set src <source prefix>
  next
  edit <sequence number>
    set gateway <IP address>
    set src <source prefix>
  next
end
```

Packet forwarding

The following sections include information about configuring packet forwarding in transparent mode:

- [MAC learning and L2 forwarding table](#)
- [Broadcast, multicast, and unicast forwarding](#)
- [Multicast processing](#)
- [Source MAC addresses](#)
- [ARP table](#)
- [Verifying the forwarding database](#)
- [STP forwarding](#)
- [Non-IPv4 Ethernet frames forwarding](#)

MAC learning and L2 forwarding table

When operating in transparent mode, a FortiGate behaves like an L2 switch in accordance with 802.1d principles:

- The forwarding database (FDB) is populated with the network devices MAC addresses during a MAC learning process, based on the source addresses seen in the Ethernet frames ingressing a FortiGate port. Static MAC

entries can also be configured using the following CLI command:

```
config system mac-address-table
edit 00:01:02:03:04:05
set interface "port3"
next
end
```

The FDB table can be verified with the following command: `diagnose netlink brctl name host TP.b`

- Ethernet IP frames forwarding is based on known MAC address on each port.
- As Spanning Tree is not running on the FortiGate, a port that comes up goes immediately into forwarding or flooding state. This last state will not occur once unicast MAC addresses are present in the FDB.



If the FortiGate in transparent mode bridges traffic to a router or host using a virtual MAC for one direction and a different physical MAC for the other direction (for example, when VRRP or HSRP protocols are used), it is highly recommended to create a static MAC entry for the virtual MAC. This is to make sure that the virtual MAC address is present in the FDB.

Broadcast, multicast, and unicast forwarding

In transparent mode, IPv4 packets are typically only forwarded by the FortiGate from a port to another port when a firewall policy is matched with action ACCEPT.

Below are exceptions.

- **L2 (IP) Broadcast frames forwarding:**



L2 (IP) means a L2 frame type 0x0800 (IP) or 0x0806 (ARP)

- **ARP:** by default, ARP broadcasts and ARP reply packets are flooded/forwarded on all ports or VLANs belonging to the same forwarding domain, without the need of firewall policies between the ports. This default behavior is necessary to allow the population of the FDB and allow further firewall policy lookup (see section transparent mode Firewall processing for more details). This option is configurable at the interface settings level with the parameter `arpforward` (enabled by default).
- **Non-ARP:** To forward non-ARP broadcasts, the following CLI command is used:

```
config system interface
edit "port2"
set broadcast-forward enable
next
end
```

- **L2 (IP) Multicast frames forwarding:** the FortiGate does not forward frames with multicast destination MAC addresses by default. Multicast traffic such as one used by routing protocols or streaming media may need to traverse the FortiGate which should not interfere this communication.

Fortinet recommends that the FortiGate is set up using Multicast policies. This allows for greater control and

predictability on traffic behavior. However Multicast traffic may be forwarded through a transparent mode device using the `multicast-skip-policy` setting. This is detailed in the section ["Multicast processing" on page 2843](#)

- **L2 (IP) Unicast frames forwarding:** a frame with a unicast destination MAC address is subject to firewall processing before being forwarded (see ["Firewall policy look up" on page 2856](#) for more details). This does not apply to ARP replies.

Multicast processing

In transparent mode, a FortiGate does not forward frames with multicast destination MAC addresses by default. If multicast traffic is required, multicast policies are recommended to allow finer control of this traffic.

Forwarding all multicast traffic with policy

Multicast traffic may have to be forwarded through a transparent mode device using the `multicast-skip-policy` system setting. This is the configuration for this solution:

```
config system settings
    set multicast-skip-policy enable
end
```

In that case, no check is performed on sources/destinations/interfaces. A multicast packet received on an interface is flooded unconditionally to all interfaces (except the incoming interface) belonging to the same forwarding domain.

Configuring firewall multicast-policy

The use of `firewall multicast-policy` allows a finer control over the multicast packets. Hereafter are some commented examples. Note that the parameter `multicast-skip-policy` mentioned above must be left to disabled.

Those policies can only be configured from the CLI.

1. Simple policy

```
config firewall multicast-policy
    edit 1
        set action accept
    next
end
```

In that case, no check is performed on sources/destinations/interfaces. A multicast packet received on an interface is flooded unconditionally to all interfaces (except the incoming interface) belonging to the same forwarding domain.

2. To restrict incoming and outgoing interfaces:

```
config firewall multicast-policy
    edit 1
        set srcintf "port1"
        set dstintf "port2"
        set action accept
    next
end
```

3. To be more restrictive (example to allow RIP2 packets from port1 to port2 and sourced by 10.10.0.10):

```
config firewall multicast-policy
edit 1
set srcintf "port1"
set srcaddr 10.10.0.10 255.255.255.255
set dstintf "port2"
set dstaddr 224.0.0.9 255.255.255.255
set action accept
next
end
```

4. This policy will allow all 224.0.0.0/255 range (OSPF, RIPv2, DVMRP...) from port1 to port2

```
config firewall multicast-policy
edit 1
set srcintf "port1"
set dstintf "port2"
set dstaddr 224.0.0.0 255.255.255
set action accept
next
end
```

Source MAC addresses

When a FortiGate is in transparent mode, it does not typically alter the original source and destination address of packets that flow through the unit. Because of this, end devices do not “see” the MAC address of the FortiGate. However, if network address translation (NAT) is enabled by a firewall policy, the source MAC address will be the MAC address of the FortiGate's management interface.

IP packets that are initiated by the FortiGate (remote management, access to FortiGuard server...) are sent in L2 Ethernet frames that have a source MAC address of the interface in the virtual domain (VDM) with the lowest MAC address. Below is an example with port2 and port3 in the same VDM, remote access done via port3, but the sniffer trace showing MAC address of port2. The address of port2 is shown in bold.

```
diagnose hardware deviceinfo nic port2
[...]
Current_HWaddr      00:09:0F:85:3F:C4
Permanent_HWaddr    00:09:0F:85:3F:C4

fgt300 (global) # diagnose hardware deviceinfo nic port3
[...]
Current_HWaddr      00:09:0F:85:3F:C5
Permanent_HWaddr    00:09:0F:85:3F:C5

diagnose sniffer packet port3 "port 80" 6
3.774236 port3 -- 192.168.171.165.2619 -> 192.168.182.136.80: syn 3961770249
0x0000  0009 0f85 3fc4 0009 0f09 3204 0800 4500      ....? .... 2...E.
0x0010  0030 8071 4000 7e06 98d7 c0a8 aba5 c0a8      .0.q@.~ .....
0x0020  b688 0a3b 0050 ec23 d109 0000 0000 7002      ...;.P.# ..... p.
0x0030  ffff d7e7 0000 0204 05b4 0101 0402
```

ARP table

In transparent mode, the Address Resolution Protocol (ARP) table is used in the following situations:

- For IP traffic received or originated by the FortiGate itself, and in destination of the management device or next-hop.
- When IPsec is used, the FortiGate uses its ARP table to forward the traffic from the IPsec tunnel to the local destination host(s).

All other forwarding decision is based on the Forwarding Database (FDB) table or optional settings.

Verifying the forwarding database

To view all instances of the forwarding database (FDB), use the following CLI command:

```
diagnose netlink brctl list
```

Example

```
FGT # diagnose netlink brctl list

list bridge information
1. root.b      fdb: size=256    used=6    num=7    depth=2    simple=no
2. mgmt.b      fdb: size=256    used=5    num=4    depth=2    simple=no
Total 2 bridges
```

Here above we can see two bridge instances for 2 VDOMs in transparent mode: `root` and `mgmt`.

- This command will dump the L2 forwarding table for each VDOM bridge instance:

```
diagnose netlink brctl name host <VDOM_name>.b
```

Example for the root VDOM:

```
FGT# diag netlink brctl name host root.b

show bridge control interface root.b host.
fdb: size=256, used=6, num=7, depth=2, simple=no
Bridge root.b host table
```

port	no	device	devname	mac addr	tvl	atributes
2	7	wan2	02:09:0f:78:69:00	0	Local Static	
5	6	trunk_1	02:09:0f:78:69:01	0	Local Static	
3	8	dmz	02:09:0f:78:69:01	0	Local Static	
4	9	internal	02:09:0f:78:69:02	0	Local Static	
3	8	dmz	00:80:c8:39:87:5a	194		
4	9	internal	02:09:0f:78:67:68	8		
1	3	wan1	00:09:0f:78:69:fe	0	Local Static	

STP forwarding

By default, the FortiGate does not forward Spanning Tree Protocol (STP) bridge protocol data units (BPDUs). If you require STP forwarding, in most configurations you must enable STP forwarding on the interface. If your FortiGate is in one-armed sniffer mode, you must instead enable STP forwarding mode on the interface.



Layer 2 loops may occur if STP is broken due to the FortiGate blocking STP BPDUs.

STP forwarding on the interface

To enable STP forwarding on an interface, use the following CLI command:

```
config system interface
  edit <interface_name>
    set stpforward enable
  next
end
```

STP forwarding mode

STP forwarding mode is an additional option that you use for STP handling when your FortiGate in one-armed sniffer mode. To configure STP forwarding mode, use the following CLI command:

```
config system interface
  edit <interface_name>
    set stpforward-mode {rpl-all-ext-id | rpl-bridge-ext-id | rpl-nothing}
  next
end
```

Non-IPv4 Ethernet frames forwarding

In the situation where non IP frames (or non Ethernet II) frames need to be accepted on a port, the parameter `l2forward` can be enabled (disabled by default). This can be used to forward frames such as PPPoE PADI, Appletalk, on other ports belonging to the same forwarding domain.

The procedure is the following:

```
config system interface
  edit port1
    set l2forward enable
  next
  edit port2
    set l2forward enable
  next
end
```

Network address translation (NAT)

While NAT is generally not used by a FortiGate in transparent mode, both [source network address translation \(SNAT\)](#) and [destination network address translation \(DNAT\)](#) can be configured.

Configuring SNAT

Source Network Address Translation (SNAT) is an option available in transparent mode and configurable in CLI only, using the following commands:

```
config firewall ippool
```

```
edit "nat-out"
    set endip 192.168.183.48
    set startip 192.168.183.48
    set interface vlan18_p3
next
end

config firewall policy
    edit 3
        set srcintf "vlan160_p2"
        set dstintf "vlan18_p3"
        set srcaddr "all"
        set dstaddr "all"
        set action accept
        set ippool enable
        set poolname "nat-out"
        set schedule "always"
        set service "ANY"
        set nat enable
    next
end
```

The sniffer trace below shows the source IP 192.168.182.93 being source translated to 192.168.183.48:

```
fgt300 (TP) # diagnose sniffer packet any "host 10.2.2.1" 4

interfaces=[any]
filters=[host 10.2.2.1]
4.891970 vlan160_p2 in 192.168.182.93 -> 10.2.2.1: icmp: echo request
4.892003 vlan18_p3 out 192.168.183.48 -> 10.2.2.1: icmp: echo request
4.892007 port3 out 192.168.183.48 -> 10.2.2.1: icmp: echo request
4.933216 vlan18_p3 in 10.2.2.1 -> 192.168.183.48: icmp: echo reply
4.933249 vlan160_p2 out 10.2.2.1 -> 192.168.182.93: icmp: echo reply
4.933253 port2 out 10.2.2.1 -> 192.168.182.93: icmp: echo reply
```

Configuring DNAT

The following example shows how to configure Destination Network Address Translation (DNAT) using a virtual IP on a FortiGate in Transparent Mode:

```
config firewall vip
    edit "vip1"
        set extip 192.168.183.48
        set extintf "vlan160_p2"
        set mappedip 192.168.182.78
    next
end

config firewall policy
    edit 4
        set srcintf "vlan160_p2"
        set dstintf "vlan18_p3"
        set srcaddr "all"
        set dstaddr "vip1"
        set action accept
        set schedule "always"
        set service "ANY"
    next
```

end



If the `mappedip` is on a different subnet than the management IP, the FortiGate must have a valid route to this destination

The sniffer trace below shows the destination IP 192.168.183.48 being translated to 192.168.182.78:

```
fgt300 (TP) # diagnose sniffer packet any "icmp" 4

interfaces=[any]
filters=[icmp]
4.126138 vlan160_p2 in 192.168.182.93 -> 192.168.183.48: icmp: echo request
4.126190 vlan18_p3 out 192.168.182.93 -> 192.168.182.78: icmp: echo request
4.126196 port3 out 192.168.182.93 -> 192.168.182.78: icmp: echo request
4.126628 vlan18_p3 in 192.168.182.78 -> 192.168.182.93: icmp: echo reply
4.126661 vlan160_p2 out 192.168.183.48 -> 192.168.182.93: icmp: echo reply
4.126667 port2 out 192.168.183.48 -> 192.168.182.93: icmp: echo reply
```

VLANs and forwarding domains

The following sections include information about configuring virtual local area networks (VLANs) and forwarding domains in transparent mode:

- [VLANs in transparent mode](#)
- [Forwarding domains in transparent mode](#)
- [VLANs vs forwarding domains](#)
- [VLAN forwarding](#)
- [Unknown VLANs and VLAN forwarding](#)
- [VLAN trunking and MAC address learning](#)
- [VLAN translation](#)

VLANs in transparent mode

A VLAN configured on a physical port is used to classify a packet in a broadcast domain in ingress and to tag packet in egress. A VLAN on the FortiGate conforms to the standard 802.1q. The following rules apply to VLAN configuration:

- a VLAN ID can be used only once on the same physical port
- the same VLAN ID can be used on a different port
- the VLAN ID range is from 1 to 4094

Forwarding domains in transparent mode

A forwarding domain is used to create separate broadcast domains and confine traffic across two or more ports. It also allows learning the same MAC in different VLANs (IVL).

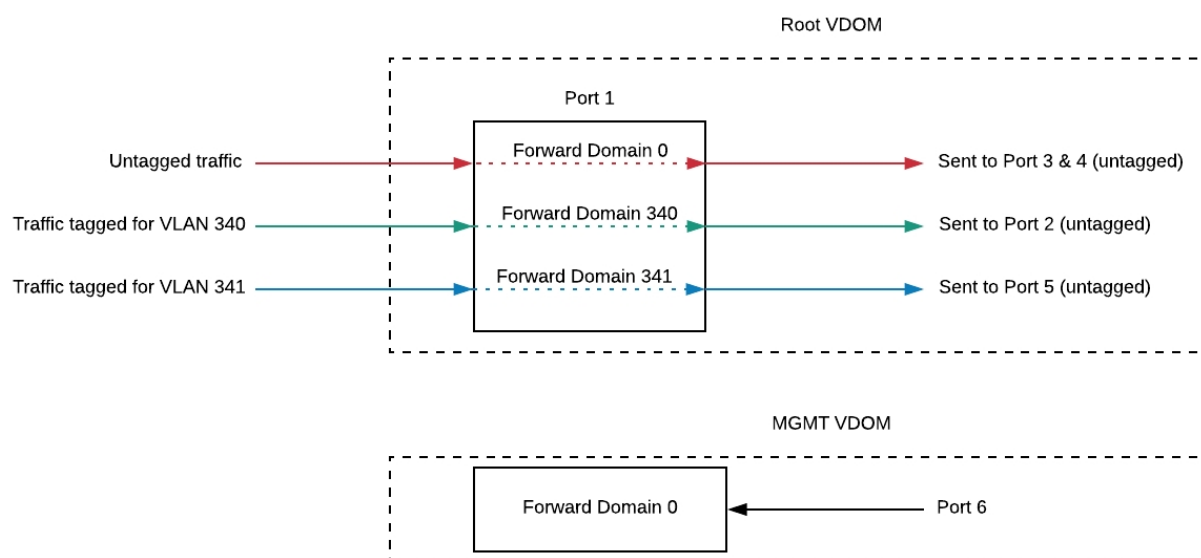
A forwarding domain and its associated ID number are unique across one VDOM, or a FortiGate with VDOMs disabled. Each new VDOM will create a new bridge instance in the FortiGate.



Even though the forwarding domain ID is not in relation with the actual VLAN numbers, it is recommended, for maintenance and troubleshooting purposes, to configure one forwarding domain per VLAN and use the same forwarding domain ID as the VLANs ID.

Once forwarding domains are configured, it is possible to configure firewall policies only between ports or VLAN belonging to the same forwarding domain.

Example configuration



This example has three forwarding domains and VLANs configured. In this example, there are two VDOMs in transparent Mode: root and MGMT. Forwarding domain 0 is the default on the FortiGate or VDOM in transparent Mode.

- Root VDOM has:
 - 3 forwarding domains, 0, 340, and 341.
 - VLAN 340 configured on port1; packets will be tagged with ID 340
 - VLAN 341 configured on port1; packets will be tagged with ID 341
 - All other ports are untagged
- MGMT VDOM has got only the default forwarding domain 0

The expected behavior is the following:

- Packets untagged ingressing port1, port3 and port4 belong to the same broadcast domain in the root VDOM
- Packets tagged with VLAN 340 ingressing port1 and Packets untagged ingressing port2 belong to the same broadcast domain in the root VDOM
- Packets tagged with VLAN 341 ingressing port1 and Packets untagged ingressing port5 belong to the same

broadcast domain in the root VDOM

- Packets untagged ingressing port6 belong to a different broadcast domain in the MGMT VDOM

CLI syntax for forwarding domain 340

```
config system interface
  edit "VLAN340"
    set forward-domain 340
    set interface "port1"
    set vlanid 340
  next
  edit "port3"
    set forward-domain 340
  next
end
```

VLANs vs forwarding domains

There are several differences between VLAN and a forwarding domain configured on a FortiGate in transparent mode:

- A forwarding domain is used to create separated broadcast domains between VLANs and allow independent VLAN learning - IVL (MAC addresses in the FDB). This would be equivalent to creating VLANs on a regular L2 switch.



When VLANs are used in the network, configuring different forwarding domains is essential to avoid broadcast duplications. See also section Default VLAN forwarding behavior for additional information.

- VLANs configured on interfaces are only used for tagging packets egressing the port and classifying packets at ingress.
- The packets processed by the direct interface (or port) itself are always sent untagged and must be received untagged.

VLAN forwarding

VLAN forwarding allows you to forward all VLANs traffic of a trunk that was connecting two network devices and where the FortiGate has been introduced, without having to perform any further configuration.

It is recommended to configure forwarding domains for each VLAN and disable this parameter in order to avoid packet from looping into the trunk from one VLAN to another. By default, the parameter `vlanforward` is disabled on each physical interface of a FortiGate or VDOM in transparent mode.

Unknown VLANs and VLAN forwarding

When a FortiGate receives a tagged frame with an unknown VLAN ID, traffic can be handled one of two ways, depending on whether VLAN forwarding is enabled or disabled.

By default, VLAN forwarding is disabled and any frames that are tagged with an unknown VLAN ID are dropped by the FortiGate.

If you enable VLAN forwarding, frames tagged with an unknown VLAN ID are forwarded from the port that received the frames to all other ports in the same forwarding domain(s). This allows you to insert the FortiGate between two devices using trunk ports without any further configuration.

VLAN trunking and MAC address learning

A FortiGate port becomes a trunk when 2 or more VLANs are configured on this port, in the same or different forwarding domains.



When trunks are configured on a FortiGate, it is essential to create forward domains, in order to avoid packets looping back on the VLANs of the trunk. This will confine all broadcasts and multicast traffic between the interfaces belonging to a same forward domain.

In the case where a trunk port is configured with a VLAN in a different forwarding domains, the MAC address of the network device connected to this port learns the FDB of each forwarding domain. This is Independent VLAN Learning (IVL).

VLAN translation

The same forwarding domain can include several different VLANs. Therefore, a frame ingressing an interface with a certain VLAN ID can be forwarded to another port with another VLAN ID. This is sometimes referred as VLAN translation.

Inter-VDOM links between NAT/route and transparent VDOMs

Inter-VDOM links between NAT/route and transparent mode VDOMs can be useful for configurations where the NAT/route VDOMs that share a common Internet service route, which can be routed through a transparent VDOM that provides additional functionality, like common Security inspection, WAN optimization, explicit proxying and so on.

Other examples include:

- Performing SSL offloading in the transparent mode VDOM and providing Internet access through a NAT/route mode VDOM.
- Applying WAN optimization in a transparent mode VDOM and other security features in the NAT/route mode VDOM.
- Using a dedicated transparent mode VDOM for the explicit web proxy in front of a NAT/route mode VDOM that applies other security features.
- An ISP configuration with multiple per-tenant NAT/route mode VDOMs all sharing a single Internet connection but where the ISP only presents a single routed subnet. Each tenant can then be assigned an IP from the subnet for their respective VDOM link interface while using a single physical port to connect to the ISP router.

For more information about inter-VDOM links, please refer to the Virtual Domains handbook.

Replay traffic scenario

Situations can arise where an identical TCP packet enters twice the FortiGate via 2 different ports. This can be due to a firewall or other network device redirecting packets out on the same port it has received it.

The FortiGate will in this condition detect a replay packet and drop it.

If the network topology or culprit devices cannot be changed to avoid this, the workaround on the FortiGate can be to disable TCP replay verification packets.

```
config system global
    set anti-replay | loose | strict | disable |
```

end

The debug flow diagnosis output hereafter shows the message indicating this condition:

```
id=20085 trace_id=179 msg="vd-VDOM_VLAN1 received a packet(proto=6, 10.10.253.9:10709
>10.10.248.5:25) from TO_EXTERNAL ."
id=20085 trace_id=179 msg="Find an existing session, id-00041475, original direction"
id=20085 trace_id=179 msg="replay packet, drop"
```

For additional diagnosis and troubleshooting procedures, go to <http://kb.fortinet.com>.

Packet forwarding using Cisco protocols

In order to pass Cisco Discover Protocol (CDP) or Cisco VLAN Trunk Protocol (VTP) packets through a FortiGate in transparent mode, the parameter `stpforward` must be applied on the port configuration. VTP and CDP packets are sent to the destination MAC address 01-00-0C-CC-CC-CC.



A Cisco NATIVE VLAN carries CDP/VTP frames. The frames of this VLAN must be received on the FortiGate physical interfaces (not VLAN sub-interface). Physical interfaces are the only ones that can send/accept non-tagged packets.

The example below will allow CDP and VTP packets to be sent from port3 up to the Remote unit, through two VDOMs, via one physical port and three port aggregations.

Port and port aggregation configuration:

```
config system interface
  edit "port1"
    set vdom "VD1"
  next
  edit "port2"
    set vdom "VD1"
  next
  edit "port3"
    set vdom "VD1"
    set stpforward enable
  next
  edit "port5"
    set vdom "VD3"
  next
  edit "port6"
    set vdom "VD3"
  next
  edit "port17"
    set vdom "VD2"
  next
  edit "port18"
    set vdom "VD2"
  next
  edit "port19"
    set vdom "VD2"
  next
  edit "port20"
    set vdom "VD2"
  next
```

```

edit "LACP_VD2_IN"
    set vdom "VD2"
    set stpforward enable
    set type aggregate
    set member "port17" "port18"
next
edit "LACP_VD2_OUT"
    set vdom "VD2"
    set stpforward enable
    set type aggregate
    set member "port19" "port20"
next
edit "LACP_VD1"
    set vdom "VD1"
    set stpforward enable
    set type aggregate
    set member "port1" "port2"
next
end

```



When using aggregation, the `stpforward` setting needs to be applied only on the port aggregation level, not on the physical port

This will also forward regular Spanning Tree BPDUs

Verification with a sniffer trace:

```
FGT# diagnose sniffer packet any "" 4
```

```

41.365434 port3 in llc unnumbered, ui, flags [command], length 72
41.365437 LACP_VD1 out llc unnumbered, ui, flags [command], length 72
41.365439 port2 out llc unnumbered, ui, flags [command], length 72
41.365479 LACP_VD2_IN in llc unnumbered, ui, flags [command], length 72
41.365482 LACP_VD2_OUT out llc unnumbered, ui, flags [command], length 72
41.365484 port19 out llc unnumbered, ui, flags [command], length 72

```


See above the CDP packet flow from port3, LACP_VD1 (port2), LACP_VD2_IN, LACP_VD2_OUT (port19).



The following sniffer trace command will filter only CDP or VTP packets :

```
FGT# diagnose sniffer packet port_name "ether host 01-00-0C-CC-CC-CC"
```

Configuration example



Step 1: Create VLANs and forwarding domains

```
config system interface
  edit "vlan102_intern"
    set forward-domain 102
    set interface "port2"
    set vlanid 102
  next
  edit "vlan102_extern"
    set forward-domain 102
    set interface "port3"
    set vlanid 102
  next
  edit "vlan103_intern"
    set forward-domain 103
    set interface "port2"
    set vlanid 103
  next
  edit "vlan103_extern"
    set forward-domain 103
    set interface "port3"
    set vlanid 103
  next
end
```

Step 2: Create the appropriate firewall policies

```
config firewall policy
  edit 1
    set srcintf "vlan102_extern"
    set dstintf "vlan102_intern"
    set srcaddr "all"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY" next
  edit 2
    set srcintf "vlan102_intern"
    set dstintf "vlan102_extern"
    set srcaddr "all"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY" next
  edit 3
    set srcintf "vlan103_intern"
    set dstintf "vlan103_extern"
    set srcaddr "all"
    set dstaddr "all"
    set action accept
    set schedule "always"
    set service "ANY" next
  edit 4
    set srcintf "vlan103_extern"
    set dstintf "vlan103_intern"
    set srcaddr "all"
    set dstaddr "all"
```

```
    set action accept
    set schedule "always"
    set service "ANY" next
end
```

Firewalls and security in transparent mode

This section contains information about using firewalls and security scanning in transparent mode. It contains the following topics:

- [Firewall policy look up](#)
- [Firewall session list](#)
- [Security scanning](#)

Firewall policy look up

In transparent mode, like in NAT/route mode, a firewall policy look up is based on the source and destination interfaces. The matching firewall policy will tell which actions to apply to the traffic, including logging and security scanning.

The FortiGate proceeds as follows to look for a matching firewall policy in transparent mode:

- **Step 1:** an Ethernet IP frame ingresses a port (or a VLAN on a port), corresponding to a specific bridge instance (from the port VDOM and Forwarding domain). This frame contains a destination MAC address that we will call MAC_D.
- **Step 2:** The FortiGate is making a MAC_D address lookup in the bridge instance to determine the port where MAC_D has been learned. This will be the destination interface.
- **Step 3:** The FortiGate is then looking for a firewall policy corresponding to the couple < source interface + destination interface >. If multiple policies with the same couple < source interface + destination interface > exist, the FortiGate screens all of them from TOP to BOTTOM (as displayed in the configuration), until a match is found. It is important to make sure that the most specific firewall policies are located at the top of the policy list, to make sure that traffic is matched to the appropriate policy.

Firewall session list

The flag **br** in the state line will indicate that this is a “bridged” session. See example below :

```
FGT# diagnose sys session list

session info: proto=17 proto_state=00 duration=59 expire=128 timeout=0 flags=000
00000 sockflag=00000000 sockport=0 av_idx=0 use=4
origin-shaper=
reply-shaper=
per_ip_shaper=
ha_id=0 hakey=0
policy_dir=0 tunnel=/
state=may_dirty br rem
statistic(bytes/packets/allow_err): org=385/5/1 reply=0/0/0 tuples=2
orgin->sink: org pre->post, reply pre->post dev=3->4/4->3 gwy=0.0.0.0/0.0.0.0
hook=pre dir=org act=noop 192.168.182.93:1025->4.2.2.1:53(0.0.0.0:0)
hook=post dir=reply act=noop 4.2.2.1:53->192.168.182.93:1025(0.0.0.0:0)
misc=0 policy_id=1 auth_info=0 chk_client_info=0 vd=1 serial=000006d3 tos=ff/ff
imp2p=0 app=0
dd_type=0 dd_rule_id=0
```

Security scanning

Security scanning occurs in the same manner in NAT/route mode and transparent mode. When a protection profile is enabled on a firewall policy for content inspection, the FortiGate acts like a transparent proxy for the protocols that need to be inspected.

The FortiGate will therefore intercept the TCP sessions and create its own session from client to server and server to client. The source and destination MAC addresses of the original L2 frames are however not altered in this communication, as described in the section Network operation : source MAC addresses in frames sent by or through the FortiGate.



Devices in the network communicating through the FortiGate do not know the presence of the FortiGate.

For more information about security scanning, see the *Security Profiles Handbook*.

IPsec VPN in transparent mode

This section contains information about configuring IPsec virtual private networks (VPNs) in transparent mode. It contains the following topics:

- [Using IPsec VPNs in transparent mode](#)
- [Example 1: Remote sites with different subnets](#)
- [Example 2: Remote sites on the same subnet](#)

Using IPsec VPNs in transparent mode

In transparent mode, IPsec VPN is supported in Policy-based configuration mode only.

IPsec VPN in transparent mode can be used in those scenarios:

- Encrypt data over routed networks without changing anything on the routers. See example 1.
- Encrypt data over a non-routed transport network (extension of a LAN for example). See example 2.

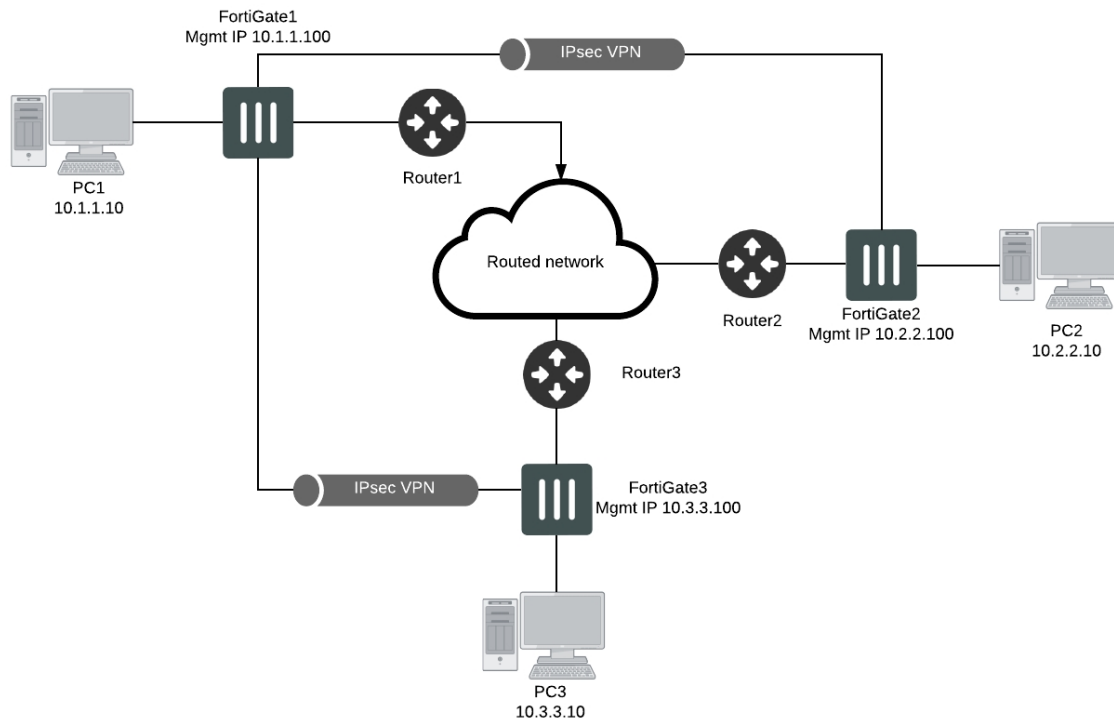
The following rules apply to IPsec in transparent mode:

- If both remote FortiGate IPsec gateways are not in the same broadcast domain (separated by routers):
 - The hosts on each side must be on different subnets.
 - The FortiGate management IP addresses must be in the same subnet as the local hosts. This is the preferred option.
- If both remote FortiGate IPsec gateways are in the same broadcast domain (separated by optical switches for examples), the hosts on each side can be :
 - On the same subnet
 - On different subnet if the appropriate static route is configured on the remote FortiGate
 - The FortiGate management IP addresses can be in any different subnet than the local hosts
- A firewall Policy with the action IPsec is used to send traffic to the remote device into the tunnel.
Therefore, it is important to place all remote devices on the appropriate ports of the FortiGate to allow a proper match < source interface + destination interface > . See section transparent mode Firewall processing for more details.



This scenario requires that the remote hosts located on the remote FortiGate's protected subnets have their MAC addresses hard coded in FortiGate's static MAC entry list. If this is not configured then it is expected to see outage in network communications.

Example 1: Remote sites with different subnets



This example provides a configuration example for IPsec VPN tunnels between three FortiGate in transparent Mode in different subnets, as well as some troubleshooting steps.

The expectation for this example is that PC1 will be able to communicate via the IPsec tunnels with PC2 and PC3, which are in different subnets.

The requirements for this example are:

- Because both FortiGate are not in the same broadcast domain (separated by routers), the hosts on each side must be on different subnets.
- FortiGate management IP addresses must be in the same subnet as the local hosts
- The default gateways (router1 ,router2, router3) for PC1 , PC2 and PC3 must be behind port2 in order for the FortiGate to match the appropriate Encrypt firewall policy (port1 --> port2)

Configuration of FortiGate 1 (FGT1):

Only relevant parts of configuration are provided.

```
config system settings
  set opmode transparent
  set manageip 10.1.1.100/255.255.255.0
end

config router static
  edit 1
    set gateway 10.1.1.254
```

```
    next
end

config firewall address
    edit "10.1.1.0/24"
        set subnet 10.1.1.0 255.255.255.0
    next
    edit "10.2.2.0/24"
        set subnet 10.2.2.0 255.255.255.0
    next
    edit "10.3.3.0/24"
        set subnet 10.3.3.0 255.255.255.0
    next
end

config vpn ipsec phase1
    edit "to_FGT2"
        set proposal 3des-sha1 aes128-sha1 des-md5
        set remote-gw 10.2.2.100
        set psksecret fortinet
    next
    edit "to_FGT3"
        set proposal 3des-sha1 aes128-sha1 des-md5
        set remote-gw 10.3.3.100
        set psksecret fortinet
    next
end

config vpn ipsec phase2
    edit "to_FGT2"
        set keepalive enable
        set phasename "to_FGT2"
        set proposal 3des-sha1 aes128-sha1
        set dst-subnet 10.2.2.0 255.255.255.0
        set src-subnet 10.1.1.0 255.255.255.0
    next
    edit "to_FGT3"
        set keepalive enable
        set phasename "to_FGT3"
        set proposal 3des-sha1 aes128-sha1
        set dst-subnet 10.3.3.0 255.255.255.0
        set src-subnet 10.1.1.0 255.255.255.0
    next
end

config firewall policy
    edit 1
        set srcintf "port1"
        set dstintf "port2"
        set srcaddr "10.1.1.0/24"
        set dstaddr "10.2.2.0/24"
        set action ipsec
        set schedule "always"
        set service "ANY"
        set inbound enable
        set outbound enable
        set vptunnel "to_FGT2"
```

```
next
edit 3
    set srcintf "port1"
    set dstintf "port2"
    set srcaddr "10.1.1.0/24"
    set dstaddr "10.3.3.0/24"
    set action ipsec
    set schedule "always"
    set service "ANY"
    set inbound enable
    set outbound enable
    set vpntunnel "to_FGT3"
next
end
```

Configuration of FortiGate 2 (FGT2):

Only relevant parts of configuration are provided.

```
config system settings
    set opmode transparent
    set manageip 10.2.2.100/255.255.255.0
end

config router static
    edit 1
        set gateway 10.2.2.254
    next
end

config firewall address
    edit "10.1.1.0/24"
        set subnet 10.1.1.0 255.255.255.0
    next
    edit "10.2.2.0/24"
        set subnet 10.2.2.0 255.255.255.0
    next
end

config vpn ipsec phase1
    edit "to_FGT1"
        set nattraversal disable
        set proposal 3des-sha1 aes128-sha1 des-md5
        set remote-gw 10.1.1.100
        set psksecret fortinet
    next
end

config vpn ipsec phase2
    edit "to_FGT1"
        set keepalive enable
        set phasename "to_FGT1"
        set proposal 3des-sha1 aes128-sha1
        set dst-subnet 10.1.1.0 255.255.255.0
        set src-subnet 10.2.2.0 255.255.255.0
    next
end
```

```

config firewall policy
  edit 1
    set srcintf "port1"
    set dstintf "port2"
    set srcaddr "10.2.2.0/24"
    set dstaddr "10.1.1.0/24"
    set action ipsec
    set schedule "always"
    set service "ANY"
    set inbound enable
    set outbound enable
    set vpntunnel "to_FGT1"
  next
end

```

Troubleshooting procedure

All steps given when PC1 pings PC2.

Verify if IPsec tunnels are up

```

FGT1 # diagnose vpn tunnel list

list all ipsec tunnel in vd 0
-----
name=to_FGT2 ver=0 serial=1 10.1.1.100:0->10.2.2.100:0 lgwy=dyn tun=tunnel mode= auto
bound_if=0
proxyid_num=1 child_num=0 refcnt=7 ilast=0 olast=0
stat: rxp=0 txp=0 rxb=0 txb=0
dpd: mode=active on=1 idle=5000ms retry=3 count=0 seqno=1455
natt: mode=none draft=0 interval=0 remote_port=0
proxyid=to_FGT2 proto=0 sa=0 ref=1 auto_negotiate=0 serial=5
  src: 10.1.1.0/255.255.255.0:0
  dst: 10.2.2.0/255.255.255.0:0

```

The above tunnel is down (output given as example)!

```

FGT2 # diagnose vpn tunnel list

list all ipsec tunnel in vd 0
-----
name=to_FGT1 10.2.2.100:0->10.1.1.100:0 lgwy=dyn tun=tunnel mode=auto bound_if=0
proxyid_num=1 child_num=0 refcnt=7 ilast=1 olast=1
stat: rxp=0 txp=0 rxb=0 txb=0
dpd: mode=active on=1 idle=5000ms retry=3 count=0 seqno=21 natt:
mode=none draft=0 interval=0 remote_port=0
proxyid=to_FGT1 proto=0 sa=1 ref=2 auto_negotiate=0 serial=1
  src: 10.2.2.0/255.255.255.0:0
  dst: 10.1.1.0/255.255.255.0:0
SA: ref=3 options=00000009 type=00 soft=0 mtu=1436 expire=1771 replaywin=0 seq no=1
life: type=01 bytes=0/0 timeout=1773/1800
dec: spi=c1a8e951 esp=3des key=24 9213fdf22b150e01abb3535d1a647044eebf772b92f2f7ee
ah=sha1 key=20 66a38bf99f0b2d234f64b5a05187995c4f56f6bb
enc: spi=322067b4 esp=3des key=24 720e5680329937fb3630b7ed70bd41bb3114d3c269ae8b61
ah=sha1 key=20 e316113eb6ea03b014b3a5f9c1a3bd386637801a

```

The above tunnel is up!

Verify that destination local hosts are seen in the ARP table (necessary for IPsec despite being in TP mode)

```
FGT2 # get system arp
```

Address	Age(min)	Hardware Addr	Interface
10.2.2.10	2	00:50:56:00:76:04	root.b
10.2.2.254	0	00:09:0f:30:29:e4	root.b

Using the debug flow command on the initiator side (example on FortiGate1)

```
FGT1 # diagnose debug flow filter addr 10.1.1.10
FGT1 # diagnose debug flow show console enable
FGT1 # diagnose debug enable
FGT1 # diagnose debug flow trace start 50

FGT1 # id=36870 trace_id=615 msg="vd-root received a packet(proto=1, 10.1.1.10:512->10.2.2.10:8) from port1."
id=36870 trace_id=615 msg="allocate a new session-00000636"
id=36870 trace_id=615 msg="Allowed by Policy-1: encrypt"
id=36870 trace_id=615 msg="enter IPsec tunnel-to_FGT2"
id=36870 trace_id=615 msg="SA is not ready yet, drop"
id=36870 trace_id=616 msg="vd-root received a packet(proto=1, 10.1.1.10:512->10.2.2.10:8) from port1."
id=36870 trace_id=616 msg="Find an existing session, id-00000636, original direction"
id=36870 trace_id=616 msg="enter IPsec tunnel-to_FGT2"
id=36870 trace_id=616 msg="encrypted, and send to 10.2.2.100 with source 10.1.1.100"
id=36870 trace_id=616 msg="send out via dev-port2, dst-mac-00:09:0f:30:29:e0"
id=36870 trace_id=617 msg="vd-root received a packet(proto=1, 10.1.1.10:512->10.2.2.10:8) from port1."
id=36870 trace_id=617 msg="Find an existing session, id-00000636, original direction"
id=36870 trace_id=617 msg="enter IPsec tunnel-to_FGT2"
id=36870 trace_id=617 msg="encrypted, and send to 10.2.2.100 with source 10.1.1.100"
id=36870 trace_id=617 msg="send out via dev-port2, dst-mac-00:09:0f:30:29:e0"
id=36870 trace_id=618 msg="vd-root received a packet(proto=1, 10.2.2.10:512->10.1.1.10:0) from port2."
id=36870 trace_id=618 msg="Find an existing session, id-00000636, reply direction"
id=36870 trace_id=618 msg="send out via dev-port1, dst-mac-00:50:56:00:76:03"
id=36870 trace_id=619 msg="vd-root received a packet(proto=1, 10.1.1.10:512->10.2.2.10:8) from port1."
id=36870 trace_id=619 msg="Find an existing session, id-00000636, original direction"
id=36870 trace_id=619 msg="enter IPsec tunnel-to_FGT2"
id=36870 trace_id=619 msg="encrypted, and send to 10.2.2.100 with source 10.1.1.100"
id=36870 trace_id=619 msg="send out via dev-port2, dst-mac-00:09:0f:30:29:e0"
id=36870 trace_id=620 msg="vd-root received a packet(proto=1, 10.2.2.10:512->10.1.1.10:0) from port2."
id=36870 trace_id=620 msg="Find an existing session, id-00000636, reply direction"
id=36870 trace_id=620 msg="send out via dev-port1, dst-mac-00:50:56:00:76:03"
```



The message "id=36870 trace_id=615 msg="SA is not ready yet, drop" simply means that the tunnel was not up yet.

Using the debug flow command on the receiver side (example on FortiGate2)

```

FGT2 # diagnose debug flow filter addr 10.1.1.10
FGT2 # diagnose debug flow show console enable
FGT2 # diagnose debug enable
FGT2 # diagnose debug flow trace start 50

FGT2 # id=36870 trace_id=51 msg="vd-root received a packet(proto=1, 10.1.1.10:512->10.2.2.10:8) from port2."
id=36870 trace_id=51 msg="allocate a new session-00000435"
id=36870 trace_id=51 msg="Allowed by Policy-1:"
id=36870 trace_id=51 msg="send out via dev-port1, dst-mac-00:50:56:00:76:04"
id=36870 trace_id=52 msg="vd-root received a packet(proto=1, 10.2.2.10:512->10.1.1.10:0) from port1."
id=36870 trace_id=52 msg="Find an existing session, id-00000435, reply direction"
id=36870 trace_id=52 msg="enter IPsec tunnel-to_FGT1"
id=36870 trace_id=52 msg="encrypted, and send to 10.1.1.100 with source 10.2.2.100"
id=36870 trace_id=52 msg="send out via dev-port2, dst-mac-00:09:0f:30:29:e4"

```

Using the sniffer trace (example on FortiGate2)

```

FGT2 # diagnose sniffer packet any "host 10.2.2.10" 4

9.460021 root.b out arp who-has 10.2.2.10 tell 10.2.2.100
9.460028 port2 out arp who-has 10.2.2.10 tell 10.2.2.100
9.460034 port1 out arp who-has 10.2.2.10 tell 10.2.2.100
9.460462 port1 in arp reply 10.2.2.10 is-at 0:50:56:0:76:4
9.460462 root.b in arp reply 10.2.2.10 is-at 0:50:56:0:76:4
[...]
49.477368 port2 in 10.1.1.10 -> 10.2.2.10: icmp: echo request
49.477444 port1 out 10.1.1.10 -> 10.2.2.10: icmp: echo request
49.477898 port1 in 10.2.2.10 -> 10.1.1.10: icmp: echo reply
50.510023 port2 in 10.1.1.10 -> 10.2.2.10: icmp: echo request
50.510079 port1 out 10.1.1.10 -> 10.2.2.10: icmp: echo request
50.510524 port1 in 10.2.2.10 -> 10.1.1.10: icmp: echo reply

```



The above ARP process in transparent mode with IPsec is allowing the FortiGate to:

- Identify the MAC address of the destination device 10.2.2.10
- Populate the MAC table (see below), which in turn will give a destination interface and allow a Firewall policy look-up

Check the FDB entries for the destination

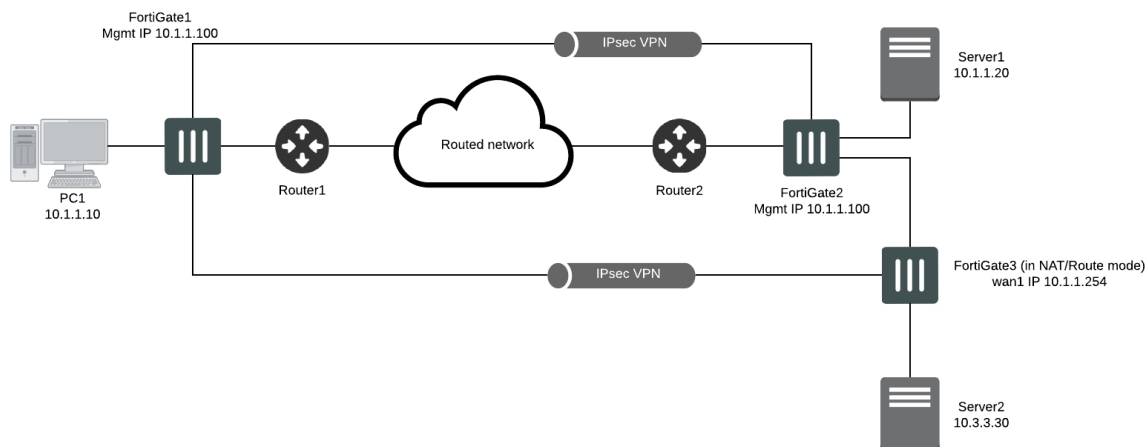
```

FGT2 # diagnose netlink brctl name host root.b

port no   device  devname  mac addr          ttl
[...]
1         2        port1    00:50:56:00:76:04  0

```

Example 2: Remote sites on the same subnet



This example provides a configuration example for IPsec VPN tunnels between two FortiGate in transparent Mode in the same subnet separated by a L2 transparent network and one remote subnet on the second site.



This scenario requires that PC1's MAC address is added to the FortiGate static MAC table. The preferred scenario would be to have a router installed between the FortiGate devices.

The expectation for this example is that PC1 will be able to communicate via the IPsec tunnel with Server1 in the same subnet, and Server2 in a different subnet.

The requirements for this example are:

- The default gateway (FGT3) for PC1 and all remote device must be behind port2 of FGT1, in order for this FortiGate to match the appropriate Encrypt firewall policy (port1 --> port2)
- Despite being in transparent mode, **FGT2 must have a valid route to Server2**
- FGT3 is used as a router between subnet 10.1.1.0/24 and 10.3.3.0/24.

PC1 MAC address added to FGT2 static MAC entries.

Server1 MAC address added to FGT1 static MAC entries.

Configuration of FortiGate 1 (FGT1):

Only relevant parts of configuration are provided.

```
config system settings
    set opmode transparent
    set manageip 10.1.1.100/255.255.255.0
end

config router static
    edit 1
        set gateway 10.1.1.252
    next
end
```



```
config system mac-address-table
  edit 00:50:56:00:76:04 ==>Server1
    set interface port2
  next
end

config firewall address
  edit "all"
  next
  edit "Server1"
    set subnet 10.1.1.20 255.255.255.255
  next
  edit "Server2"
    set subnet 10.3.3.30 255.255.255.255
  next
  edit "10.1.1.0/24"
    set subnet 10.1.1.0 255.255.255.0
  next
  edit "gateway"
    set subnet 10.1.1.254 255.255.255.255
  next
end

config vpn ipsec phase1
  edit "to_FGT2"
    set proposal 3des-sha1 aes128-sha1 des-md5
    set remote-gw 10.1.1.200
    set psksecret fortinet
  next
end

config vpn ipsec phase2
  edit "to_FGT2"
    set keepalive enable
    set phasename "to_FGT2"
    set proposal 3des-sha1 aes128-sha1
    set src-subnet 10.1.1.0 255.255.255.0
  next
end

config firewall policy
  edit 1
    set srcintf "port1"
    set dstintf "port2"
    set srcaddr "10.1.1.0/24"
    set dstaddr "Server1"
    set action ipsec
    set schedule "always"
    set service "ANY"
    set inbound enable
    set outbound enable
    set vpntunnel "to_FGT2"
  next
  edit 2
    set srcintf "port1"
    set dstintf "port2"
    set srcaddr "10.1.1.0/24"
```

```
        set dstaddr "Server2"
        set action ipsec
        set schedule "always"
        set service "ANY"
        set inbound enable
        set outbound enable
        set vpntunnel "to_FGT2"
    next
edit 3
    set srcintf "port1"
    set dstintf "port2"
    set srcaddr "10.1.1.0/24"
    set dstaddr "gateway"
    set action ipsec
    set schedule "always"
    set service "ANY"
    set inbound enable
    set outbound enable
    set vpntunnel "to_FGT2"
next
end
```



Firewall Policy 3 is not mandatory and is only used to allow PC1 to test a ping reachability to its default gateway 10.1.1.254.

Configuration of FortiGate 2 (FGT2):

Only relevant parts of configuration are provided.

```
config system settings
    set opmode transparent
    set manageip 10.1.1.200/255.255.255.0
end

config router static
    edit 1
        set gateway 10.1.1.252
    next
    edit 2
        set dst 10.3.3.0 255.255.255.0
        set gateway 10.1.1.254
    next
end

config system mac-address-table
    edit 00:50:56:00:76:03
        set interface wan1
    next
end

config firewall address
    edit "all"
    next
    edit "PC1"
```

```
        set subnet 10.1.1.10 255.255.255.255
    next
    edit "10.1.1.0/24"
        set subnet 10.1.1.0 255.255.255.0
    next
    edit "10.3.3.0/24"
        set subnet 10.3.3.0 255.255.255.0
    next
end

config vpn ipsec phase1
    edit "to_FGT1"
        set proposal 3des-sha1 aes128-sha1 des-md5
        set remote-gw 10.1.1.100
        set psksecret fortinet
    next
end

config vpn ipsec phase2
    edit "to_FGT1"
        set keepalive enable
        set phaselname "to_FGT1"
        set proposal 3des-sha1 aes128-sha1
        set dst-subnet 10.1.1.0 255.255.255.0
    next
end

config firewall policy
    edit 1
        set srcintf "internal"
        set dstintf "wan1"
        set srcaddr "10.1.1.0/24" set dstaddr "PC1"
        set action ipsec
        set schedule "always"
        set service "ANY" set inbound enable
        set outbound enable
        set vpntunnel "to_FGT1"
    next
    edit 2
        set srcintf "internal"
        set dstintf "wan1"
        set srcaddr "10.3.3.0/24" set dstaddr "PC1"
        set action ipsec
        set schedule "always"
        set service "ANY" set inbound enable
        set outbound enable
        set vpntunnel "to_FGT1"
    next
end
```

Troubleshooting procedure

Check the ARP entries of PC1

```
C:\ arp -a
```

```
Interface: 10.1.1.10 --- 0x20003
```

Internet Address	Physical Address	Type
10.1.1.20	00-50-56-00-76-04	dynamic
10.1.1.254	00-09-0f-85-3f-c8	dynamic



MAC address **00-09-0f-85-3f-c8** is the FGT3 interface in subnet 10.1.1.0/24.

FDB entries of FGT1

```
FGT1 (global) # diagnose netlink brctl name host Vdom1.b
```

```
show bridge control interface Vdom1.b host. fdb:
size=256, used=6, num=6, depth=1
Bridge Vdom1.b host table
```

port no	device	devname	mac addr	tvl	attributes
1	10	port1	00:50:56:00:76:03	0	
2	9	port2	00:50:56:00:76:04	44	static
2	9	port2	00:09:0f:85:3f:c8	13	
1	10	port1	00:09:0f:88:2f:69	0	Local Static
2	9	port2	00:09:0f:88:2f:68	0	Local Static
2	9	port2	00:09:0f:23:01:d6	0	



MAC address **00:09:0f:23:01:d6** is “internal” port MAC address of FGT2 00:09:0F:23:01:D6. This is the MAC address used for management in the transparent mode VDOM of FGT2, chosen between the lowest MAC address between wan1 (00:09:0F:78:00:74) and internal (00:09:0F:23:01:D6).

ARP entries of FGT2

```
FGT2 (TP) # get system arp
```

Address	Age(min)	Hardware Addr	Interface
10.1.1.20	82	00:50:56:00:76:04	TP.b
10.1.1.100	13	00:09:0f:88:2f:68	TP.b
10.1.1.254	76	00:09:0f:85:3f:c8	TP.b



it is important to have the entry for 10.1.1.254 which is the route to 10.3.3.0/24 .

IPsec Tunnel verification on FGT1

```
FGT1 (Vdom1) # diagnose vpn tunnel list
```

```
list all ipsec tunnel in vd 3
```

```
-----
name=to_FGT2 10.1.1.100:0->10.1.1.200:0 lgwy=dyn tun=tunnel mode=auto bound_if=0
proxyid_num=1 child_num=0 refcnt=10 ilast=0 olast=0
stat: rxp=2754 txp=2945 rxb=308448 txb=176700
dpd: mode=active on=1 idle=5000ms retry=3 count=0 seqno=166 natt:
```

```

mode=none draft=0 interval=0 remote_port=0
proxyid=to_FGT2 proto=0 sa=1 ref=2 auto_negotiate=0 serial=1
src: 10.1.1.0/255.255.255.0:0
dst: 0.0.0.0/0.0.0.0:0
SA: ref=3 options=00000009 type=00 soft=0 mtu=1436 expire=1271 replaywin=0 seqno=1e1
life:type=01 bytes=0/0 timeout=1750/1800
dec: spi=3f148cb7 esp=3des key=24 834832201a0dbbf60b0098106f08380538dbd94cacdlad31
ah=sha1 key=20 b0257a135cba745b956bef3d4b8a6e65934c074b
enc: spi=1895305e esp=3des key=24 4d3092f0b3f84184d4779f85a9953230bf9bc28bd93c0afa
ah=sha1 key=20 0c70acf6ad2193ec5934e2a4332fd09f32016e60
npu_flag=00 npu_rgw=10.1.1.200 npu_lgw=10.1.1.100 npu_selid=0

```

Sniffer trace on FGT1 when PC1 pings all 3 remote destinations

```

FGT1 (Vdom1) # diagnose sniffer packet any "icmp" 4

interfaces=[any]
filters=[icmp]
0.342268 port1 in 10.1.1.10 -> 10.3.3.30: icmp: echo request
0.342844 port2 in 10.3.3.30 -> 10.1.1.10: icmp: echo reply
0.342884 port1 out 10.3.3.30 -> 10.1.1.10: icmp: echo reply
0.771700 port1 in 10.1.1.10 -> 10.1.1.20: icmp: echo request
0.772504 port2 in 10.1.1.20 -> 10.1.1.10: icmp: echo reply
0.772539 port1 out 10.1.1.20 -> 10.1.1.10: icmp: echo reply
0.907377 port1 in 10.1.1.10 -> 10.1.1.254: icmp: echo request
0.907850 port2 in 10.1.1.254 -> 10.1.1.10: icmp: echo reply
0.907883 port1 out 10.1.1.254 -> 10.1.1.10: icmp: echo reply

```

Sniffer trace on FGT1 filtered on IPsec protocol

```

FGT1 (Vdom1) # diagnose sniffer packet port2 "proto 50" 6

interfaces=[port2]
filters=[proto 50]
pcap_lookupnet: port2: no IPv4 address assigned

1.249003 port2 -- 10.1.1.100 -> 10.1.1.200: ip-proto-50 92
0x0000 0009 0f23 01d6 0009 0f88 2f68 0800 4500 ...# ..... /h..E.
0x0010 0070 c9e6 0000 3f32 9a48 0a01 0164 0a01 .p ... ?2.H...d..
0x0020 01c8 1895 305f 0000 01e2 02b6 37b6 8b2c ....0_ ..... 7...,

1.249478 port2 -- 10.1.1.200 -> 10.1.1.100: ip-proto-50 92
0x0000 0009 0f88 2f68 0009 0f23 01d6 0800 4500 ..../h...# ... E.
0x0010 0070 2e31 0000 3f32 35fe 0a01 01c8 0a01 .p.1..?25 .....
0x0020 0164 3f14 8cb8 0000 01e2 324d 66e2 9236 .d? ..... 2Mf..6

```



From the above trace, the MAC address **0009 0f88 2f68** is the MAC address of FGT1 port2 . This is the MAC address used for management in the transparent mode VDOM of FGT1, chosen between the lowest MAC address between port1 (00:09:0F:88:2F:69) and port2 ((00:09:0F:88:2F:68).

Debug flow on FGT1 filtered on Server3

```

FGT1 (Vdom1) # diagnose debug flow filter addr 10.3.3.30
FGT1 (Vdom1) # diagnose debug flow show console enable

```

```
FGT1 (Vdom1) # diagnose debug enable
FGT1 (Vdom1) # diagnose debug flow trace start 10

id=20085 trace_id=11 msg="vd-Vdom1 received a packet(proto=1, 10.1.1.10:512->10.3.3.30:8)
    from port1."
id=20085 trace_id=11 msg="Find an existing session, id-00004e85, original direction"
id=20085 trace_id=11 msg="enter IPsec tunnel-to_FGT2"
id=20085 trace_id=11 msg="encrypted, and send to 10.1.1.200 with source 10.1.1.100"
id=20085 trace_id=11 msg="send out via dev-port2, dst-mac-00:09:0f:23:01:d6"
id=20085 trace_id=12 msg="vd-Vdom1 received a packet(proto=1, 10.3.3.30:512->10.1.1.10:0)
    from port2."
id=20085 trace_id=12 msg="Find an existing session, id-00004e85, reply direction"
id=20085 trace_id=12 msg="send out via dev-port1, dst-mac-00:50:56:00:76:03"
```



From the trace above, **dst-mac-00:09:0f:23:01:d6** is "internal" port MAC address of FGT2 00:09:0F:23:01:D6. This is the MAC address used for management in the transparent mode VDOM of FGT2, chosen between the lowest MAC address between wan1 (00:09:0F:78:00:74) and internal (00:09:0F:23:01:D6).

Using FortiManager and FortiAnalyzer

FortiManager and FortiAnalyzer are supported similarly to NAT/route mode. For more information about this please consult the Fortinet documentation at <http://docs.fortinet.com> or the Knowledge base at <http://kb.fortinet.com>.

Establishing a communication to a FortiAnalyzer is done as per the example hereafter (from global level if VDOM is enabled). This setting is independent from being in transparent mode. However, as stated earlier in this section the management VDOM must have IP connectivity to the FortiAnalyzer.

```
FGT (global) # show system fortianalyzer

config system fortianalyzer
    set status enable
    set server 10.2.2.2
end
```

High availability in transparent mode

This section contains information about configuring high availability in transparent mode. It contains the following topics:

- [Virtual clustering](#)
- [MAC address assignment](#)



For complete information about HA, please refer to the FortiGate Administration Guide or the HA Technical guides available at <http://docs.fortinet.com> or the Knowledge Base at <http://kb.fortinet.com>.

Any other statement and feature description in this document apply to a FortiGate Cluster running in Active-Passive mode.

Virtual clustering

If VDOM (virtual domain) is enabled on a cluster operating transparent mode, HA virtual clustering can be configured in active-passive mode.

This will provide:

- Failover protection between two instances of a VDOM operating on two different FortiGate in the cluster.
- Load balancing between the FortiGate devices on a per-VDOM basis.

The roles have been defined such as, in normal operation:

- FortiGate1 is Master for VDOM1 and Slave for VDOM2
- FortiGate2 is Master for VDOM2 and Slave for VDOM1

In case of a failure or reboot of a FortiGate, the remaining unit will become Master for VDOM1 and VDOM2.



The VDOMs given in this example are showing physical ports but a VDOM can also include VLAN interfaces.



The L2 connectivity between the FortiGate is showing 4 separate L2 switches, but it could also be one single switch on each side configured with appropriate VLANs.

Configuration example

- FortiGate1:

```
FGT1 (global) # show system ha

config system ha
    set mode a-p
    set hbdev "port5" 0 "port6" 0
```



```
set vcluster2 enable
set override disable
set priority 200
    config secondary-vcluster
        set override enable
        set priority 100
        set vdom "VDOM2"
    end
end
```

- **FortiGate2:**

```
FGT2 (global) # show system ha

config system ha
    set mode a-p
    set hbdev "port5" 0 "port6" 0
    set vcluster2 enable
    set override disable
    set priority 200
        config secondary-vcluster
            set override enable
            set priority 100
            set vdom "VDOM2"
        end
end
```

MAC address assignment

If a cluster is operating in transparent mode, the FortiGate Clustering Protocol (FGCP) assigns a virtual MAC address for the Master unit management IP address. Since you can connect to the management IP address from any interface, all of the FortiGate interfaces appear to have the same virtual MAC address.

Best practices

1. Create forwarding domains when VLANs are used and set `vlanforward` to `disable` on all relevant physical interface.
2. The forward-domain ID can be different to the VLAN ID, but it is recommended for troubleshooting and readability to keep them the same.
3. Only interfaces from the same forwarding domains can have firewall policies between each others.
4. In order to allow IVL (independent VLAN learning), the VLANs must be placed in separate forwarding domains.
5. If an out-of-band management is desired, use if possible a VDOM in NAT/route mode as management VDOM and create (an) other transparent mode VDOM(s) for the user traffic.
6. As Spanning Tree BPDUs are not forwarded by default, insert the FortiGate with caution to avoid L2 loops.
7. Multicast packets are not forwarded by default; this might cause routing protocols (RIP2, OSPF) disruption.
8. When using HSRP or VRRP configure static MAC entries for the Virtual MAC addresses.

Chapter 27 - Troubleshooting

This handbook chapter presents troubleshooting and problem solving issues that may help you with your FortiGate and contains the following sections:

- ["Troubleshooting methodologies" on page 2877](#) walks you through best practice concepts of FortiOS troubleshooting.
- ["Troubleshooting tools" on page 2881](#) describes some of the basic commands and parts of FortiOS that can help you with troubleshooting.
- ["Troubleshooting tips" on page 2908](#) presents most of the common issues and how to address them.
- ["Troubleshooting resources" on page 2936](#) identifies Fortinet's resources for troubleshooting.

Troubleshooting methodologies

This section explains how to prepare for troubleshooting, create a troubleshooting plan, and where to find additional resources.

The following topics are covered:

- [Ensure you have administrator-level access to required equipment](#)
- [Establish a baseline](#)
- [Define the problem](#)
- [Create a troubleshooting plan](#)
- [Obtain any required equipment](#)
- [Consult Fortinet resources](#)

Ensure you have administrator-level access to required equipment

Before troubleshooting your FortiGate, you will need administrator access to the equipment. If you're a client on a FortiGate that has virtual domains (VDOMs) enabled, you can often troubleshoot within your own VDOM. However, you should inform the super admin for the FortiGate that you'll be performing troubleshooting tasks.

Also, you may need access to other networking equipment, such as switches, routers, and servers, to carry out tests. If you don't have access to this equipment, contact your network administrator for assistance.

Establish a baseline

A FortiGate operates at all layers of the OSI model. For this reason, troubleshooting problems can become complex. Establishing baseline parameters for your system before a problem occurs helps to reduce the complexity when you need to troubleshoot.

Many of the guiding questions in the following sections serve to compare the current problem situation to normal operation on your FortiGate. A best practice is to establish and record the normal operating status. Regular operation data shows trends, and allows you to see when changes occur and when there may be a problem. You can gather this data by using logs and SNMP tools to monitor the system performance or by regularly running information gathering commands and saving the output.



Back up your FortiOS configuration on a regular basis. This is a good practice and not only for troubleshooting. You can restore the backed up configuration as needed and save the time and effort of recreating it from the factory default settings.

You can use the following CLI commands to obtain normal operating data for a FortiGate:

```
get system status
```

Displays firmware versions and FortiGuard engine versions, and other system information

<code>get system performance status</code>	Displays CPU and memory states, average network usage, average sessions and session setup rate, virus caught, IPS attacks blocked, and uptime
<code>get hardware memory</code>	Displays information about memory
<code>get system session status</code>	Displays total number of sessions
<code>get router info routing-table all</code>	Displays all the routes in the routing table, including their type, source, and other useful data
<code>get ips session</code>	Displays memory used and maximum amount available to IPS as well and counts
<code>get webfilter ftgd-statistics</code>	Displays a list of FortiGuard related counts of status, errors, and other data
<code>diagnose system session list</code>	Displays the list of current detailed sessions
<code>show system dns</code>	Displays the configured DNS servers
<code>diagnose system ntp status</code>	Displays information about NTP servers

These commands are just a sample. You can run any commands for information gathering that apply to your system. For example, if you have active VPN connections, use the `get vpn *` series of commands to get more information about them.

To see an extensive snapshot of your system, you can use the `execute tac report` command. This command runs many diagnostic commands for specific configurations. Regardless of the features deployed on your FortiGate, this command records the current state of each feature. Then, if you need to perform troubleshooting later, you can run the same command again and compare the differences to quickly identify any suspicious output.

Define the problem

The following questions help you define the problem. Be as specific as possible with your answers. Once you define the problem, you can search for a solution and then create a plan for how to solve it.

- **What is the problem?**

The problem being observed is not necessarily the actual problem. You should determine where the problem lies before starting to troubleshoot the FortiGate.

- **Was the device working before?**

If the device never worked, it might be defective. For more information, see [Troubleshooting your FortiGate Installation](#) in the [Getting Started](#) chapter.

- **Can you reproduce the problem ?**

If the problem is intermittent, it may be dependent on system load. Note that it may be difficult to troubleshoot an intermittent problem because it's difficult to reproduce.

- **What has changed?**

Don't assume that nothing has changed in the network. Use the FortiGate event log to identify any possible configuration changes. There may be changes in the operating environment. For example, there might be a gradual increase in load as more sites are forwarded through the firewall.

If something has changed, roll back the change and assess the impact.

- **What is the scope of the problem?**

After you isolate the problem, determine what applications, users, devices, and operating systems the problem affects.

- What's not working? Be specific.
- Is there more than one thing that isn't working?
- Is it partly working? If so, what parts are working?
- Is it a connectivity issue for the entire device, or is there an application that isn't reaching the Internet?
- Where did the problem occur?
- When did the problem occur and to which users or groups of users?
- What components are involved?
- What applications are affected?
- Can you use a packet sniffer to trace the problem?
- Can you use system debugging or look in the session table to trace the problem?
- Do any of the log files indicate a failure has occurred?

The answers to these questions help you narrow down the problem and identify what you should check during your troubleshooting. The more things you can eliminate, the fewer things you need to check during troubleshooting. For this reason, be as specific and accurate as you can when you gather information.

Create a troubleshooting plan

Once you define the problem and gather facts, you can create a troubleshooting plan to solve the problem.

You should list all possible causes of the problem and how you can test for each cause.

The plan acts as a checklist so that you know what you've tried and what's left to check. This is also important to have if more than one person is performing troubleshooting tasks.

Be ready to add to your plan, as needed.

Providing supporting elements

If you contact Fortinet's Technology Assistance Center (TAC), be prepared to provide the following information:

- Firmware build version (use the `get system status` command)
- Network topology diagram
- Recent configuration file
- Recent debug log (optional)
- Summary of troubleshooting steps that you've already taken and the results.



Don't provide the output from the `exec tac` report unless TAC requests it. The output from this command is very large and isn't required in many cases.

Obtain any required equipment

To test your solution, you may require additional networking equipment, computers, or other equipment.

Network administrators usually have additional networking equipment available to loan you, or a lab where you can bring the FortiGate unit to test.

If you don't have access to equipment, check for shareware applications that can perform the same tasks. Often, there are software solutions that you can use when hardware is too expensive.

Consult Fortinet resources

After you define your problem, create a plan to find a solution, and carry out that plan. If you can't resolve the problem, see ["Troubleshooting resources" on page 2936](#).

Troubleshooting tools

FortiOS provides several tools that can help you troubleshoot both hardware and software issues. These tools include diagnostics and ports. You use ports when you need to understand the traffic coming in and going out on a specific port. For example, the FortiGate uses port UDP 53 for DNS and RBL lookups.

This section contains the following topics:

- [FortiOS diagnostics](#)
- [FortiOS ports](#)
- [FortiAnalyzer/FortiManager ports](#)
- [FortiGuard troubleshooting](#)

FortiOS diagnostics

FortiOS has a collection of diagnostic commands that you can use to troubleshoot and monitor the performance of your network. The `get` and `diagnose` CLI commands are the two main groups of diagnostic commands. Both commands display information about system resources, connections, and settings that allow you to locate problems or monitor system performance.

This topic includes diagnostics commands to help with:

- [Date and time](#)
- [Resource usage](#)
- [Proxy operation](#)
- [Hardware NIC](#)
- [Traffic trace](#)
- [Session table](#)
- [Firewall session setup rate](#)
- [Finding object dependencies](#)
- [Flow trace](#)
- [Packet sniffing and packet capture](#)
- [NPU-based interfaces](#)
- [Debug command](#)
- [The execute tac report command](#)
- [Other commands](#)

Date and time

The system date and time are important for FortiGuard services, logging events, and sending alerts. The wrong time makes the log entries confusing and difficult to use.

Use Network Time Protocol (NTP) to set the date and time, if possible. This is an automatic method that doesn't require manual intervention. However, you must ensure that the port is allowed through the firewalls on your network. FortiToken synchronization requires NTP in many situations.

How to set the date and time - GUI

1. Go to the **System Information** widget on the **Dashboard**. The date and time are displayed next to **System Time**.
2. To adjust the date and time settings, go to **System > Settings**.
3. In **System Time**, you can set the time zone, date and time, and select **NTP usage**.

How to check the date and time - CLI

You can check the date and time using the CLI commands `execute date` and `execute time`.

```
config system global
    set timezone <integer>
end
config system ntp
    set type custom
    config ntpserver
        edit 1
            set server "ntp1.fortinet.net"
        next
        edit 2
            set server "ntp2.fortinet.net"
        next
    end
set ntpsync enable
set syncinterval 60
end
```

Use the `set timezone ?` command to display a list of timezones and the integers that represent them.

Resource usage

Each program that runs on a computer has one or more processes associated with it. For example, if you open a Telnet program, it has an associated Telnet process. The same is true in FortiOS. All processes share the system resources in FortiOS, including memory and CPU.

Use the `get system performance status` command to show the FortiOS performance status.

Sample output:

```
FGT# get system performance status
CPU states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
CPU0 states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
CPU1 states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
CPU2 states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
CPU3 states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
Memory: 4050332k total, 527148k used (13%), 3381312k free (83%), 141872k freeable (3%)
Average network usage: 41 / 28 kbps in 1 minute, 54 / 44 kbps in 10 minutes, 42 / 34
kbps in 30 minutes
Average sessions: 33 sessions in 1 minute, 48 sessions in 10 minutes, 38 sessions in 30
minutes
Average session setup rate: 0 sessions per second in last 1 minute, 0 sessions per
second in last 10 minutes, 0 sessions per second in last 30 minutes
Virus caught: 0 total in 1 minute
IPS attacks blocked: 0 total in 1 minute
```

Uptime: 0 days, 22 hours, 59 minutes

Monitor the CPU and memory usage of internal processes, using the following command:

```
get system performance top <delay> <max_lines>
```

The data that the command lists includes the name of the daemon, the process ID, whether the process is sleeping or running, the CPU use percentage, and the memory use percentage.

Sample output:

```
get system performance top 10 100
Run Time: 0 days, 23 hours and 4 minutes
OU, ON, OS, 100I, OWA, OHI, OSI, OST; 3955T, 3298F
httpsd 212 S 0.4 0.6
forticron 169 S 0.4 0.4
newcli 4054 R 0.4 0.2
reportd 174 S 0.0 1.4
pyfcgid 325 S 0.0 0.8
cmdbsvr 141 S 0.0 0.7
miglogd 160 S 0.0 0.6
httpsd 211 S 0.0 0.6
src-vis 180 S 0.0 0.6
pyfcgid 327 S 0.0 0.6
pyfcgid 328 S 0.0 0.6
pyfcgid 329 S 0.0 0.6
httpsd 162 S 0.0 0.5
cw_acd 189 S 0.0 0.5
httpsd 3998 S 0.0 0.5
httpsd 4050 S 0.0 0.5
updated 176 S 0.0 0.5
httpsd 4052 S 0.0 0.4
miglogd 203 S 0.0 0.4
miglogd 204 S 0.0 0.4
```

Proxy operation

Monitor proxy operations, using the following command:

```
diagnose test application <application> <option>
```

To display a list of available <application> values, enter:

```
diagnose test application ?
```

The <option> value depends on the application value that you use in the command. To display a list of available <option> values, enter:

```
diagnose test application <application> ?
```

For example, if the application is http, the CLI command that displays the <option> values is:

```
diagnose test application http ?
```

Hardware NIC

To monitor hardware network operations, use the following command:

```
diagnose hardware deviceinfo nic <interface>
```

The information that this command shows is important because errors at the interface indicate data link or physical layer issues which may impact the performance of the FortiGate.

The following example shows a sample output when you set <interface> to lan:

```
System_Device_Name lan
Current_HWaddr 00:09:0f:68:35:60
Permanent_HWaddr 00:09:0f:68:35:60
State up
Link up
Speed 100
Duplex full
[.....]
Rx_Packets=5685708
Tx_Packets=4107073
Rx_Bytes=617908014
Tx_Bytes=1269751248
Rx_Errors=0
Tx_Errors=0
Rx_Dropped=0
Tx_Dropped=0
[....]
```

The `diagnose hardware deviceinfo nic` command displays a list of error names and values that are related to hardware. The following table describes possible hardware errors:

Field	Description
Rx_Errors = rx error count	Bad frame was marked as error by PHY
Rx_CRC_Errors + Rx_Length_Errors - Rx_Align_Errors	This error is only valid in 10/100M mode
Rx_Dropped or Rx_No_Buffer_Count	Running out of buffer space
Rx_Missed_Errors	Equals Rx_FIFO_Errors + CEXTERR (Carrier Extension Error Count); only valid in 1000M mode, which is marked by PHY
Tx_Errors = Tx_Aborted_ Errors	ECOL (Excessive Collisions Count); only valid in half-duplex mode

Field	Description
	Late Collisions (LATECOL) Count
Tx_Window_Errors	Late collisions are collisions that occur after 64-byte time into the transmission of the packet while working in 10 to 100 Mb/s data rate and 512-byte time into the transmission of the packet while working in the 1,000 Mb/s data rate. This register only increments if transmits are enabled and the device is in half-duplex mode.
Rx_Dropped	See Rx_Errors
Tx_Dropped	Not defined
Collisions	Total number of collisions experienced by the transmitter; valid in half-duplex mode
Rx_Length_Errors	Transmission length error
Rx_Over_Errors	Not defined
Rx_CRC_Errors	Frame CRC error
Rx_Frame_Errors	Same as Rx_Align_Errors This error is only valid in 10/100M mode.
Rx_FIFO_Errors	Same as Rx_Missed_Errors - a missed packet count
Tx_Aborted_Errors	See Tx_Errors
Tx_Carrier_Errors	The PHY should assert the internal carrier sense signal during every transmission. Failure to do so may indicate that the link has failed or the PHY has an incorrect link configuration. This register only increments if transmits are enabled. This register isn't valid in internal SerDes 1 mode (TBI mode for the 82544GC/EI) and is valid only when the Ethernet controller is operating at full duplex.
Tx_FIFO_Errors	Not defined
Tx_Heartbeat_Errors	Not defined
Tx_Window_Errors	See LATECOL
Tx_Single_Collision_Frames	Counts the number of times that a successfully transmitted packet encountered a single collision The value increments only if transmits are enabled and the Ethernet controller is in half-duplex mode.

Field	Description
Tx_Multiple_Collision_Frames	A Multiple Collision Count which indicates the number of times that a transmit encountered more than one collision, but less than 16. The value increments only if transmits are enabled and the Ethernet controller is in half-duplex mode.
Tx_Deferred	Counts defer events A deferred event occurs when the transmitter can't immediately send a packet due to the medium being busy because another device is transmitting, the IPG timer hasn't expired, half-duplex deferral events are occurring, XOFF frames are being received, or the link isn't up. This register only increments if transmits are enabled. This counter doesn't increment for streaming transmits that are deferred due to TX IPG.
Rx_Frame_Too_Longs	The Rx frame is oversized
Rx_Frame_Too_Shots	The Rx frame is too short
Rx_Align_Errors	This error is only valid in 10/100M mode
Symbol Error Count	Counts the number of symbol errors between reads - SYMERRS. The count increases for every bad symbol that's received, whether or not a packet is currently being received and whether or not the link is up. This register increments only in internal SerDes mode.

Traffic trace

Traffic tracing allows you to follow a specific packet stream. This is useful to confirm that packets are taking the route you expect them to take on your network.

View the characteristics of a traffic session through specific security policies using:

```
diagnose system session
```

Trace per-packet operations for flow tracing using:

```
diagnose debug flow
```

Trace per-Ethernet frame using:

```
diagnose sniffer packet
```

Trace a route from a FortiGate to a destination IP address using:

```
# execute traceroute www.fortinet.com
traceroute to www.fortinet.com (66.171.121.34), 32 hops max, 84 byte packets
1 172.20.120.2 0.637 ms 0.653 ms 0.279 ms
2 209.87.254.221 <static-209-87-254-221.storm.ca> 2.448 ms 2.519 ms 2.458 ms
3 209.87.239.129 <core-2-g0-2.storm.ca> 2.917 ms 2.828 ms 9.324 ms
4 209.87.239.199 <core-3-bdi1739.storm.ca> 13.248 ms 12.401 ms 13.009 ms
5 216.66.41.113 <v502.core1.tor1.he.net> 17.181 ms 12.422 ms 12.268 ms
6 184.105.80.9 <100ge1-2.core1.nyc4.he.net> 21.355 ms 21.518 ms 21.597 ms
```

```

7 198.32.118.41 <ny-paix-gni.twgate.net> 83.297 ms 84.416 ms 83.782 ms
8 203.160.228.217 <217-228-160-203.TWGATE-IP.twgate.net> 82.579 ms 82.187 ms 82.066 ms
9 203.160.228.229 <229-228-160-203.TWGATE-IP.twgate.net> 82.055 ms 82.455 ms 81.808 ms
10 203.78.181.2 82.262 ms 81.572 ms 82.015 ms
11 203.78.186.70 83.283 ms 83.243 ms 83.293 ms
12 66.171.127.177 84.030 ms 84.229 ms 83.550 ms
13 66.171.121.34 <www.fortinet.com> 84.023 ms 83.903 ms 84.032 ms
14 66.171.121.34 <www.fortinet.com> 83.874 ms 84.084 ms 83.810 ms

```

Session table

A session is a communication channel between two devices or applications across the network. Sessions allow FortiOS to inspect and act on a sequential group of packets in a session all together instead of inspecting each packet individually. Each of these sessions has an entry in the session table that includes important information about the session.

Use as a tool

Session tables are useful troubleshooting tools because they allow you to verify open connections. For example, if you have a web browser open to browse the Fortinet website, you would expect a session entry from your computer, on port 80, to the IP address for the Fortinet website. Another troubleshooting method is if there are too many sessions for FortiOS to process, you can examine the session table for evidence why this is happening.

You can view the FortiGate session table from either the FortiGate GUI or CLI. The most useful troubleshooting data comes from the CLI. The session table in the GUI also provides useful summary information, particularly the current policy number that the session is using.

GUI session information

You can view session information by going to the **FortiView** page. Read more about [FortiView consoles](#) in the Handbook's FortiView chapter.

How to find which security policy a specific connection is using

Every program and device on your network must have a communication channel, or session, open to pass information. The FortiGate manages these sessions with features such as traffic shaping, antivirus scanning, and blocking known bad web sites. Each session has an entry in the session table.

You may want to find information for a specific session for troubleshooting. For example, if a secure web browser session isn't working properly, you can check the session table to ensure the session is still active and going to the proper address. The session table can also tell you the security policy number it matches, so you can check what's happening in that policy.

1. Know your connection information.

You need to be able to identify the session you want. For this, you need the source IP address (usually your computer), the destination IP address (if you have it), and the port number which is determined by the program that you're using. Some common ports are:

- Port 80 (HTTP for web browsing),
- Port 22 (SSH used for secure login and file transfers)
- Port 23 (Telnet for a text connection)
- Port 443 (HTTPS for secure web browsing)

2. Find your session and policy ID.

Go to **FortiView > All Sessions**. Find your session by finding your source IP address, destination IP address (if you have it), and port number. The policy ID is listed after the destination information. If the list of sessions is very long, you can filter the list to make it easier to find your session.

3. When there are many sessions, use a filter to help you find your session.

If there are multiple pages of sessions, it's difficult to find a single session. You can use a filter to block out sessions that you don't want. Click the search icon on the column heading to select the filter. Select **Source IP** and enter your source IP address. Now, only sessions that originate from your IP address are displayed in the session table. If the list is still too long, you can do the same for the **Source port**. That makes it easy to find your session and the security policy ID.

CLI session information

The session table output that the `diagnose system session list` command generates is very large. You can use filters to display only the session data that you're interested in.

An entry is placed in the session table for each traffic session passing through a security policy. The following command lists the information for a session in the table:

```
diagnose system session list
```

The filter option displays specific information, for example:

```
diagnose system session filter <option>
```

The values for `<option>` include the following:

<code>clear</code>	Clear session filter
<code>dintf</code>	Destination interface
<code>dport</code>	Destination port
<code>dst</code>	Destination IP address
<code>duration</code>	Duration of the session
<code>expire</code>	Expire
<code>negate</code>	Inverse filter
<code>nport</code>	NAT'd source port
<code>nsrc</code>	NAT'd source ip address
<code>policy</code>	Policy ID
<code>proto</code>	Protocol number
<code>proto-state</code>	Protocol state

<code>session-state1</code>	Session state1
<code>session-state2</code>	Session state2
<code>sintf</code>	Source interface
<code>sport</code>	Source port
<code>src</code>	Source IP address
<code>vd</code>	Index of virtual domain, -1 matches all

Even though UDP is a sessionless protocol, the FortiGate keeps track of the following states:

- When UDP reply doesn't have a value of 0
- When UDP reply has a value of 1

The following table displays firewall session states from the session table:

State	Description
<code>log</code>	Session is being logged
<code>local</code>	Session is originated from or destined for local stack
<code>ext</code>	Session is created by a firewall session helper
<code>may_dirty</code>	Session is created by a policy For example, the session for <code>ftp control channel</code> will have this state but <code>ftp data channel</code> won't. This is also seen when NAT is enabled.
<code>ndr</code>	Session will be checked by IPS signature
<code>nds</code>	Session will be checked by IPS anomaly
<code>br</code>	Session is being bridged (TP) mode

Firewall session setup rate

The number of sessions that can be established in a set period of time is useful information. A session is an end-to-end TCP/IP connection for communication with a limited lifespan. If you record the setup rate during normal operation, when you experience problems you can compare the baseline setup rate to the rate that occurs when you're troubleshooting. This can be a useful step to help you define your problem.

A reduced firewall session setup rate can be the result of a number of things, such as a lack of system resources on the FortiGate or reaching the limit of your session count for your VDOM.

To view your session setup rate method 1 - CLI

```
FGT# get system performance status
CPU states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
```



```

CPU0 states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
CPU1 states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
CPU2 states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
CPU3 states: 0% user 0% system 0% nice 100% idle 0% iowait 0% irq 0% softirq
Memory: 4050332k total, 530512k used (13%), 3376844k free (83%), 142976k freeable (3%)
Average network usage: 131 / 90 kbps in 1 minute, 26 / 15 kbps in 10 minutes, 49 / 42
kbps in 30 minutes
Average sessions: 80 sessions in 1 minute, 30 sessions in 10 minutes, 42 sessions in 30
minutes
Average session setup rate: 3 sessions per second in last 1 minute, 0 sessions per
second in last 10 minutes, 0 sessions per second in last 30 minutes
Virus caught: 0 total in 1 minute
IPS attacks blocked: 0 total in 1 minute
Uptime: 1 days, 2 hours, 45 minutes

```

The information you're looking for is the average sessions section, in the above output. This example shows that there were 80 sessions in 1 minute, or an average of 3 sessions per second. The values for 10 minutes and 30 minutes allow you to take a longer average for a more reliable value if your FortiGate is working at maximum capacity. The smallest FortiGate can have 1,000 sessions established per second across the unit.

Remember that session setup rate is a global command. If you have multiple VDOMs configured with many sessions in each one, the session setup rate per VDOM will be slower than if there are no VDOMs configured.

Finding object dependencies

An administrator may not be permitted to delete a configuration object if there are other configuration objects that depend on it. This command identifies other objects which depend on, or make reference to, the configuration object in question. If an error is displayed that an object is in use and can't be deleted, this command can help identify the source of the problem.

Additionally, if you have a virtual interface with objects that depend on it, you need to find and remove those dependencies before you delete the interface.

CLI method

When you run multiple VDOMs, you use this command in the global configuration only and it searches for the named object in both the most recently used global and VDOM configurations:

```
diagnose system checkused <path.object.mkey>
```

For example, to verify which objects a security policy with an ID of 1 refers to, enter the following command:

```
diagnose system checkused firewall.policy.policyid 1
```

To check what is referred to by interface `port1`, enter the following command:

```
diagnose system checkused system.interface.name port1
```

To show all dependencies for an interface, enter the following command:

```
diagnose system checkused system.interface.name <interface name>
```

Sample output:

```

entry used by table firewall.address:name '10.98.23.23_host'
entry used by table firewall.address:name 'NAS'
entry used by table firewall.address:name 'all'
entry used by table firewall.address:name 'fortinet.com'

```

```
entry used by table firewall.vip:name 'TORRENT_10.0.0.70:6883'
entry used by table firewall.policy:policyid '21'
entry used by table firewall.policy:policyid '14'
entry used by table firewall.policy:policyid '19'
```

In this example, the interface has dependent objects, including four address objects, one VIP, and three security policies.

GUI method

In the GUI, you can easily check and remove the object dependencies for an interface.

To remove interface object dependencies - GUI

1. Go to **Network > Interfaces**.
The **Ref.** column displays the number of objects that refer to this interface.
2. Select the number in the **Ref.** column for the interface.
A window listing the dependencies appears.
3. Use these detailed entries to locate and remove object references to this interface.
The trash can icon changes from gray when you've removed all object dependencies.
4. Remove the interface by selecting the check box for the interface, and select **Delete**.

Flow trace

To trace the flow of packets through the FortiGate, use the following command:

```
diagnose debug flow trace start
```

Follow packet flow by setting a flow filter, using this command:

```
diagnose debug flow {filter | filter6} <option>
```

If your network uses IPv4, enter `filter`. If your network uses IPv6, enter `filter6`.

One of the following variables replaces `<option>`:

Variable	Description
<code>addr</code>	IPv4 or IPv6 address
<code>clear</code>	clear filter
<code>daddr</code>	destination IPv4 or IPv6 address
<code>dport</code>	destination port
<code>negate</code>	inverse IPv4 or IPv6 filter
<code>port</code>	port

Variable	Description
proto	protocol number
saddr	source address
sport	source port
vd	index of virtual domain; -1 matches all



`diagnose debug flow` output is recorded as event log messages and which are then sent to a FortiCloud or a FortiAnalyzer, if connected. Don't run this command longer than necessary, since it generates significant amounts of data.

To start flow monitoring with a specific number of packets - CLI:

```
diagnose debug flow trace start <N>
```

To stop flow tracing at any time using - CLI:

```
diagnose debug flow trace stop
```

The following example shows the flow trace for a device with an IP address of 203.160.224.97:

```
diagnose debug enable
diagnose debug flow filter addr 203.160.224.97
diagnose debug flow show function-name enable
diagnose debug flow trace start 100
```

Flow trace output example - HTTP

To observe the debug flow trace, connect to the web site at the following address:

```
https://www.fortinet.com
```

Comment: SYN packet received:

```
id=20085 trace_id=209 func=resolve_ip_tuple_fast
line=2700 msg="vd-root received a packet(proto=6,
192.168.3.221:1487->203.160.224.97:80) from port5."
```

SYN sent and a new session is allocated:

```
id=20085 trace_id=209 func=resolve_ip_tuple line=2799
msg="allocate a new session-00000e90"
```

Lookup for next-hop gateway address:

```
id=20085 trace_id=209 func=vf_ip4_route_input line=1543
msg="find a route: gw-192.168.11.254 via port6"
```

Source NAT, lookup next available port:

```
id=20085 trace_id=209 func=get_new_addr line=1219
```

```
msg="find SNAT: IP-192.168.11.59, port-31925"
direction"
```

Matched security policy. Check to see which policy this session matches:

```
id=20085 trace_id=209 func=fw_forward_handler line=317
msg="Allowed by Policy-3: SNAT"
```

Apply source NAT:

```
id=20085 trace_id=209 func=__ip_session_run_tuple
line=1502 msg="SNAT 192.168.3.221->192.168.11.59:31925"
```

SYN ACK received:

```
id=20085 trace_id=210 func=resolve_ip_tuple_fast line=2700
msg="vd-root received a packet(proto=6, 203.160.224.97:80-
>192.168.11.59:31925) from port6."
```

Found existing session ID. Identified as the reply direction:

```
id=20085 trace_id=210 func=resolve_ip_tuple_fast line=2727
msg="Find an existing session, id-00000e90, reply direction"
```

Apply destination NAT to inverse source NAT action:

```
id=20085 trace_id=210 func=__ip_session_run_tuple
line=1516 msg="DNAT 192.168.11.59:31925-
>192.168.3.221:1487"
```

Lookup for next-hop gateway address for reply traffic:

```
id=20085 trace_id=210 func=vf_ip4_route_input line=1543
msg="find a route: gw-192.168.3.221 via port5"
```

ACK received:

```
id=20085 trace_id=211 func=resolve_ip_tuple_fast line=2700
msg="vd-root received a packet(proto=6,
192.168.3.221:1487->203.160.224.97:80) from port5."
```

Match existing session in the original direction:

```
id=20085 trace_id=211 func=resolve_ip_tuple_fast line=2727
msg="Find an existing session, id-00000e90, original
direction"
```

Apply source NAT:

```
id=20085 trace_id=211 func=__ip_session_run_tuple
line=1502 msg="SNAT 192.168.3.221->192.168.11.59:31925"
```

Receive data from client:

```
id=20085 trace_id=212 func=resolve_ip_tuple_fast
line=2700 msg="vd-root received a packet(proto=6,
192.168.3.221:1487->203.160.224.97:80) from port5."
```

Match existing session in the original direction:

```
id=20085 trace_id=212 func=resolve_ip_tuple_fast
line=2727 msg="Find an existing session, id-00000e90,
original direction"
```

Apply source NAT:

```
id=20085 trace_id=212 func=__ip_session_run_tuple
line=1502 msg="SNAT 192.168.3.221->192.168.11.59:31925"
```

Receive data from server:

```
id=20085 trace_id=213 func=resolve_ip_tuple_fast
line=2700 msg="vd-root received a packet(proto=6,
203.160.224.97:80->192.168.11.59:31925) from port6."
```

Match existing session in reply direction:

```
id=20085 trace_id=213 func=resolve_ip_tuple_fast
line=2727 msg="Find an existing session, id-00000e90,
reply direction"
```

Apply destination NAT to inverse source NAT action:

```
id=20085 trace_id=213 func=__ip_session_run_tuple
line=1516 msg="DNAT 192.168.11.59:31925-
>192.168.3.221:1487"
```

Flow trace output example - IPsec (policy-based)

```
id=20085 trace_id=1 msg="vd-root received a packet(proto=1, 10.72.55.240:1->10.71.55.10:8)
from internal."
id=20085 trace_id=1 msg="allocate a new session-00001cd3"
id=20085 trace_id=1 msg="find a route: gw-66.236.56.230 via wan1"
id=20085 trace_id=1 msg="Allowed by Policy-2: encrypt"
id=20085 trace_id=1 msg="enter IPsec tunnel-RemotePhase1"
id=20085 trace_id=1 msg="encrypted, and send to 15.215.225.22 with source 66.236.56.226"
id=20085 trace_id=1 msg="send to 66.236.56.230 via intf-wan1"
id=20085 trace_id=2 msg="vd-root received a packet (proto=1, 10.72.55.240:1-1071.55.10:8)
from internal."
id=20085 trace_id=2 msg="Find an existing session, id-00001cd3, original direction"
id=20085 trace_id=2 msg="enter IPsec ="encrypted, and send to 15.215.225.22 with source
66.236.56.226" tunnel-RemotePhase1"
id=20085 trace_id=2 msgid=20085 trace_id=2 msg="send to 66.236.56.230 via intf-wan1"
```

Packet sniffing and packet capture

When you troubleshoot networks, it helps to look inside the header of the packets. This helps to determine if the packets, route, and destination are all what you expect. Packet capture can also be called a network tap, packet sniffing, or logic analyzing. FortiOS devices can sniff packets using CLI commands or capture packets using the GUI.

Packet sniffing using CLI commands is well-suited for spot checking traffic, but if you have complex filters to enter it can be a lot of work to enter them each time. You can also save the sniffing output. However, you must log to a file and then analyze the file later.

Packet capture in the GUI makes it easy for you to set up multiple filters at once and run only one or two as you need to. You can also use controls to start and stop capturing when you want to. You download packet capture

output to your local computer as a *.pcap file. You must use a third party application, such as Wireshark, to read *.pcap files. This method is useful to send information to Fortinet support to help resolve an issue.

The following table presents a comparison between the two methods:

Features	Packet sniffing	Packet capture
Command location	CLI	GUI
Third party software required	puTTY to log plaintext output	Wireshark, or similar application, to read *.pcap files
Read output in plain text file	yes	no
Read output as *.pcap file using Wireshark, or similar application	no	yes
Easily configure single quick and simple filter	yes	no
Record packet interface	yes	no
Configure complex sniffer filters on multiple interface	no	yes
sniff IPv6	hard	easy
sniff non-IP packets	no	yes
Filter packets by protocol and/or port	easy	easy
Filter packets by source and/or destination address	easy	easy

Packet sniffing

If you're running a constant traffic application, such as ping, packet sniffing can tell you if the traffic is reaching the destination, what the port of entry is on the FortiGate unit, if the ARP resolution is correct, and if the traffic is being sent back to the source as expected.

Sniffing packets can also tell you if the FortiGate is silently dropping packets for reasons such as Reverse Path Forwarding (RPF). RPF, also called anti-spoofing, prevents an IP packet from being forwarded if its source IP doesn't belong to a locally attached subnet (local interface) or isn't part of the routing between the FortiGate and another source (static route, RIP, OSPF, BGP). Note that you can disable RPF by turning on asymmetric routing in the CLI (`config system settings, set asymroute enable`), but this disables stateful inspection on the FortiGate and causes many features to be turned off.



If you configure virtual IP addresses on your FortiGate, it will use those addresses instead of the physical IP addresses. You'll notice this when you're sniffing packets because all the traffic uses the virtual IP addresses. This is due to the ARP update that's sent out when the VIP address is configured.

Before you start sniffing packets on the CLI, you should prepare to capture the output to a file. A large amount of data may scroll by and you won't be able to see it without first saving it to a file. One method is to use a terminal program like puTTY to connect to the FortiGate CLI. Once the packet sniffing count is reached, you can end the session and analyze the output in the file.

The general form of the internal FortiOS packet sniffer command is:

```
diagnose sniffer packet <interface_name> <'filter'> <verbose> <count> <tsformat>
```

To stop the sniffer, type CTRL+C.

<interface_name>	The name of the interface to sniff, such as "port1" or "internal". This can also be "any" to sniff all interfaces.
<'filter'>	What to look for in the information the sniffer reads. "none" indicates no filtering, and all packets are displayed as the other arguments indicate. The filter must be inside single quotes (').
<verbose>	The level of verbosity as one of: 1 - print header of packets 2 - print header and data from IP of packets 3 - print header and data from Ethernet of packets 4 - print header of packets with interface name
<count>	The number of packets the sniffer reads before stopping. If you don't put a number here, the sniffer will run until you stop it with <CTRL+C>.
<tsformat>	The format of timestamp. <ul style="list-style-type: none"> • a: absolute UTC time, yyyy-mm-dd hh:mm:ss.ms • l: absolute LOCAL time, yyyy-mm-dd hh:mm:ss.ms • otherwise: relative to the start of sniffing, ss.ms

For a simple sniffing example, enter the CLI command `diagnose sniffer packet port1 none 1 3`. This displays the next three packets on the port1 interface using no filtering, and verbose level 1. At this verbosity level, you can see the source IP and port, the destination IP and port, action (such as ack), and sequence numbers.

In the output below, port 443 indicates these are HTTPS packets and that 172.20.120.17 is both sending and receiving traffic.

```
Head_Office_620b # diagnose sniffer packet port1 none 1 3
interfaces=[port1]
filters=[none]

0.545306 172.20.120.17.52989 -> 172.20.120.141.443: psh 3177924955 ack 1854307757

0.545963 172.20.120.141.443 -> 172.20.120.17.52989: psh 1854307757 ack 3177925808

0.562409 172.20.120.17.52988 -> 172.20.120.141.443: psh 4225311614 ack 3314279933
```

For a more advanced example of packet sniffing, the following commands will report packets on any interface that are traveling between a computer with the host name of “PC1” and a computer with the host name of “PC2”. With verbosity 4 and above, the sniffer trace displays the interface names where traffic enters or leaves the FortiGate unit. Remember to stop the sniffer, type `CTRL+C`.

```
FGT# diagnose sniffer packet any "host <PC1> or host <PC2>" 4
```

or

```
FGT# diagnose sniffer packet any "(host <PC1> or host <PC2>) and icmp" 4
```

The following CLI command for a sniffer includes the ARP protocol in the filter which may be useful to troubleshoot a failure in the ARP resolution (for example, PC2 may be down and not responding to the FortiGate ARP requests).

```
FGT# diagnose sniffer packet any "host <PC1> or host <PC2> or arp" 4
```

Packet capture

Packet capture tells you what is happening on the network at a low level. This can be very useful for troubleshooting problems, such as:

- Finding missing traffic
- Seeing if sessions are setting up properly
- Locating ARP problems such as broadcast storm sources and causes
- Confirming which address a computer is using on the network, if they have multiple addresses or are on multiple networks
- Confirming routing is working as you expect
- Connecting wireless clients
- Missing PING packets
- A particular type of packet is having problems, such as UDP, which is commonly used for streaming video

If you're running a constant traffic application such as ping, packet capture can tell you if the traffic is reaching the destination, how the port enters and exits the FortiGate, if the ARP resolution is correct, and if the traffic is returning to the source as expected. You can also use packet switching to verify that NAT or other configuration is translating addresses or routing traffic the way that you want it to.

Before you start capturing packets, you need to have a good idea of what you're looking for. Capture is used to confirm or deny your ideas about what's happening on the network. If you try capture without a plan to narrow your search, you can end up with too much data to effectively analyze. On the other hand, you need to capture enough packets to really understand all of the patterns and behavior that you're examining.

To use packet capture, the FortiGate must have a disk. You can enable the `capture-packet` in the firewall policy, using the following CLI commands:

```
config firewall policy
  edit <id>
    set capture-packet enable
  end
```

To configure packet capture filters, go to **Network > Packet Capture**.

When you add a packet capture filter, enter the following information and select **OK**.

Interface	Select the interface to sniff from the drop-down menu. You must select one interface. You can't change the interface without deleting the filter and creating a new one, unlike the other fields.
Max Packets to Save	Enter the number of packets to capture before the filter stops. This number can't be zero. You can halt the capturing before this number is reached.
Enable Filters	Select this option to specify filter fields
Host(s)	Enter the IP address of one or more hosts Separate multiple hosts with commas. To enter a range, use a dash without spaces, for example 172.16.1.5-172.16.1.15, or enter a subnet.
Port(s)	Enter one or more ports to capture on the selected interface. Separate multiple ports with commas. To enter a range, use a dash without spaces, for example 88-90
VLAN(s)	Enter one or more VLANs (if any). Separate multiple VLANs with commas.
Protocol	Enter one or more protocols. Separate multiple protocols with commas. To enter a range, use a dash without spaces, for example 1-6, 17, 21-25.
Include IPv6 Packets	Select this option if you're troubleshooting IPv6 networking, or if your network uses IPv6. Otherwise, leave it disabled.
Include Non-IP Packets	The protocols in the list are all IP based except for ICMP (ping). To capture non-IP based packets, select this feature. Examples of non-IP packets include IPsec, IGMP, ARP, and ICMP.

If you select a filter, you have the option to start and stop packet capture in the edit window, or download the captured packets. You can also see the filter status and the number of packets captured.

You can select the filter and start capturing packets. When the filter is running, the number of captured packets increases until it reaches the **Max Packet Count** or you stop it. When the filter is running, you can't download the output file.

When the packet capture is complete, you can download the *.pcap file. You must use a third party application, such as Wireshark, to read *.pcap files. This tool provides you with extensive analytics and the full contents of the packets that were captured.

To start, stop, or resume packet capture, use the symbols on the screen. These symbols are the same as those used for audio or video playback. Hover over the symbol to reveal explanatory text. Similarly, to download the *.pcap file, use the download symbol on the screen.

NPU-based interfaces

Many Fortinet products contain network processors, such as NP1, NP2, NP4, and NP6, which means that offloading requirements vary depending on the model.

When you use the NPU-based interfaces, you can see only the initial session setup will be seen using the `diagnose debug flow` command. If the session is correctly programmed into the ASIC (fastpath), the `diagnose debug flow` command won't detect the packets that arrive at the CPU. If the NPU functionality is disabled, the CPU detects all the packets. However, you should only this for troubleshooting purposes.

First, obtain the NP4 or NP6 ID and the port numbers, using the following command:

```
diagnose npu {np4|npu6}list
```

Sample output:

```
ID Model Slot Interface
0 On-board port1 fabric1 fabric3 fabric5
1 On-board fabric2 port2 base2 fabric4
```

Run the following commands:

```
diagnose npu {np4|npu6}fastpath disable <dev_id>
```

(where `dev_id` is the NP4 or NP6 number)

Then, run this command:

```
diagnose npu {np4|npu6}fastpath-sniffer enable port1
```

Sample output:

```
NP4 Fast Path Sniffer on port1 enabled
```

This causes traffic on **port1** of the network processor to be sent to the CPU. this means that you can take a standard sniffer trace and use other diagnose commands, if it's a standard CPU-driven port.

These commands only apply to the newer NP4 and NP6 interfaces.

Debug command

Debug output provides continuous, real-time event information and continues until you stop it or reboot the unit. Debug output can affect system performance and is continually generated even though output might not be displayed in the CLI console.

Debug information that's displayed in the console scrolls in the console display and may prevent you from entering CLI commands, such as, the command to disable the debug display. To turn off debug output as the display scrolls by, press the **↑** key to recall the recent `diagnose debug` command, press backspace, and type "0", and **Enter**.

To enable debug output display, use the following command:

```
diagnose debug enable
```

Once you enable debug output display, specify the debug information that you require using the following command:

```
diagnose debug <option> <level>
```

Debug command options include the following:

<code>enable</code>	Enable debug output
<code>disable</code>	Disable debug output
<code>info</code>	Show active debug level settings
<code>reset</code>	Reset all debug level to default
<code>report</code>	Report for tech support
<code>crashlog</code>	Crash log info
<code>config-error-log</code>	Configure error log info
<code>sql-log-error</code>	SQL log database error info
<code>application</code>	application
<code>kernel</code>	kernel
<code>remote-extender</code>	remote-extender
<code>cli</code>	Debug CLI
<code>cmdb-trace</code>	Trace CLI
<code>rating</code>	Display rating info
<code>authd</code>	Authentication daemon
<code>fsso-polling</code>	FSSO active directory poll module
<code>flow</code>	Trace packet flow in kernel
<code>urlfilter</code>	urlfilter
<code>admin</code>	Admin user

You can set the debug level at the end of the command. For example, typical values are 2 and 3, like in the following commands:

```
diagnose debug application DHCPDS 2
diagnose debug application spamfilter 2
```

Fortinet support will advise you about which debugging level you should use.

You can enable timestamps to the debug output, using the following command:

```
diagnose debug console timestamp enable
```

When you finish examining the debug output, disable it, using the following command:

```
diagnose debug disable
```

Debug output example

This example shows the IKE negotiation for a secure logging connection from a FortiGate to a FortiAnalyzer.

```
diagnose debug reset
diagnose vpn ike log-filter src-addr4 192.168.11.2
diagnose debug enable
```

Sample output:

```
FGh_FtiLog1: IPsec SA connect 0 192.168.11.2->192.168.10.201:500, natt_mode=0 rekey=0
phase2=FGh_FtiLog1
FGh_FtiLog1: using existing connection, dpd_fail=0
FGh_FtiLog1: found phase2 FGh_FtiLog1
FGh_FtiLog1: IPsec SA connect 0 192.168.11.2 -> 192.168.10.201:500 negotiating
FGh_FtiLog1: overriding selector 225.30.5.8 with 192.168.11.2
FGh_FtiLog1: initiator quick-mode set pfs=1536...
FGh_FtiLog1: try to negotiate with 1800 life seconds.
FGh_FtiLog1: initiate an SA with selectors: 192.168.11.2/0.0.0.0->192.168.10.201,
ports=0/0, protocol=0/0
Send IKE Packet(quick_outI1):192.168.11.2:500(if0) -> 192.168.10.201:500, len=348
Initiator: sent 192.168.10.201 quick mode message #1 (OK)
FGh_FtiLog1: set retransmit: st=168, timeout=6.
```

In this example:

192.168.11.2->192.168.10.201:500	Source and destination gateway IP address
dpd_fail=0	Found existing Phase 1
pfs=1536...	Create new Phase 2 tunnel

The execute tac report command

`exec tac report` is an execute command that runs an exhaustive series of diagnostic commands. It runs commands that are only needed if you're using certain features, such as HA, VPN tunnels, or a modem. The report takes a few minutes to finish because of the amount of output that's generated. If you're logging CLI output to a file, you can run this command to familiarize yourself with the diagnostic commands.

When you contact [Fortinet Support](#), you may be asked to use the output from this CLI command to provide information about your FortiGate and its current state.

Other commands

ARP table

To view the ARP cache, use the following command:

```
get system arp
```

To view the ARP cache in the system, use the following command:

```
diagnose ip arp list
```

Sample output:

```
index=14 ifname=internal 224.0.0.5 01:00:5e:00:00:05 state=00000040 use=72203
confirm=78203 update=72203 ref=1
index=13 ifname=dmz 192.168.3.100 state=00000020 use=1843 confirm=650179 update=644179
ref=2 ? VIP
index=13 ifname=dmz 192.168.3.109 02:09:0f:78:69:ff state=00000004 use=71743 confirm=75743
update=75743 ref=1
index=14 ifname=internal 192.168.11.56 00:1c:23:10:f8:20 state=00000004 use=10532
confirm=10532 update=12658 ref=4
```

To remove the ARP cache, use the following command:

```
execute clear system arp table
```

To remove a single ARP entry, use the following command:

```
diagnose ip arp delete <interface name> <IP address>
```

To add static ARP entries, use the following command:

```
config system arp-table
```

Time and date settings

Check time and date settings for log message timestamp synchronization (the Fortinet support group may request this) and for certificates that have a time requirement to check for validity. Use the following commands:

```
execute time
current time is: 12:40:48
last ntp sync:Thu Mar 16 12:00:21 2006
execute date
current date is: 2006-03-16
```

To force synchronization with an NTP server, use the following command:

```
config system ntp
set ntpsync {enable|disable}
end
```

If all devices have the same time, it helps to correlate log entries from different devices.

IP address

There may be times when you want to verify that the IP addresses assigned to the FortiGate interfaces are what you expect them to be. This is easily accomplished from the CLI, using the following command:

```
diagnose ip address list
```

The output from this command lists the IP address and mask (if available), the `index` of the interface (a type of ID number), and the `devname` (the interface name). While physical interface names are set, virtual interface names can vary. A good way to use this command is to list all of the virtual interface names. Listing all the virtual interface names is a good use of this command. For `vsys_ha` and `vsys_fgfm`, the IP addresses are the local host, which are virtual interfaces that are used internally.

```
# diagnose ip address list
```

```

IP=10.31.101.100->10.31.101.100/255.255.255.0 index=3 devname=internal
IP=172.20.120.122->172.20.120.122/255.255.255.0 index=5 devname=wan1
IP=127.0.0.1->127.0.0.1/255.0.0.0 index=8 devname=root
IP=127.0.0.1->127.0.0.1/255.0.0.0 index=11 devname=vsys_ha
IP=127.0.0.1->127.0.0.1/255.0.0.0 index=13 devname=vsys_fgfm

```

FortiOS ports

There are 65,535 ports in TCP and UDP stacks that applications can use when they communicate with each other. Many of these ports are commonly known to be associated with specific applications or protocols. These ports can be useful when you troubleshoot your network.

Use the following ports when you troubleshoot your FortiGate:

Port	Functionality
UDP 53	DNS lookup, RBL lookup
UDP 53 or UDP 8888	FortiGuard Antispam or Web Filtering rating lookup
UDP 53 (default) or UDP 8888 and UDP 1027 or UDP 1031	FDN server list - source and destination port numbers vary by originating or reply traffic
UDP 123	NTP synchronization
UDP 162	SNMP traps
UDP 514	SYSLOG - All FortiOS versions can use syslog to send log messages to remote syslog servers
TCP 22	Configuration backup to FortiManager unit or FortiGuard Analysis and Management Service
TCP 25	SMTP alert email, encrypted virus sample auto-submit
TCP 389 or TCP 636	LDAP or PKI authentication
TCP 443	FortiGuard Antivirus or IPS update - When you request updates from a FortiManager, instead of directly from the FDN, you must reconfigure this port as TCP 8890
TCP 443	FortiGuard Analysis and Management Service
TCP 514	FortiGuard Analysis and Management Service log transmission (OFTP)
TCP 514	SSL Management Tunnel to FortiGuard Analysis and Management Service
TCP 514	Quarantine, remote access to logs and reports on a FortiAnalyzer unit, device registration with FortiAnalyzer units (OFTP)
TCP 1812	RADIUS authentication

Port	Functionality
TCP 8000 and TCP 8002	FSSO
TCP 10151	FortiGuard Analysis and Management Service contract validation

FortiAnalyzer and FortiManager ports

If you have a FortiAnalyzer or FortiManager on your network, you may need to use the following ports to troubleshoot network traffic:

Port	Functionality
UDP 53	DNS lookup
UDP 123	NTP synchronization
UDP 137-138	Windows share
UDP 162	SNMP traps
UDP 514	Syslog, log forwarding
TCP 21 or TCP 22	Log and report upload
TCP 25	SMTP alert email
TCP 389 or TCP 636	User name LDAP queries for reports
TCP 443	RVS update
TCP 1812	RADIUS authentication
TCP 3000	Log aggregation client

FortiGuard troubleshooting

The FortiGuard service provides updates to AntiVirus (AV), Antispam (AS), Intrusion Protection Services (IPS), Webfiltering (WF), and more. The FortiGuard Distribution System (FDS) consists of a number of servers across the world that provide updates to your FortiGate unit. Problems can occur with the connection to FDS and its configuration on your local FortiGate unit. Some of the more common troubleshooting methods are listed here, including:

- [Troubleshooting process for FortiGuard updates](#)
- [FortiGuard server settings](#)

Troubleshooting process for FortiGuard updates

The following process shows the logical steps that you should take when you troubleshoot problems with FortiGuard update:

1. Does the device have a valid licence that includes these services?

Each device requires a valid FortiGuard license to access updates for some or all of these services. You can verify the status of the support contract for your devices at the [Fortinet Support](#) website.

2. If the device is part of a high availability (HA) cluster, do all members of the cluster have the same level of support?

As with the previous step, you can verify the status of the support contract for all of the devices in your HA cluster at the [Fortinet Support](#) website.

3. Are services enabled on the device?

To see the FortiGuard information and status for a device, in the GUI, go to **System > FortiGuard**. On that page, you can verify the status of each component, and enable each service. If you encounter problems, see the [FortiGuard](#) discussion in the [Fortinet Communication Ports and Protocols](#) chapter of the FortiOS Handbook.

4. Can the device communicate with FortiGuard servers?

Go to System > FortiGuard in the GUI and try to update AV and IPS, or test the availability of WF and AS default and alternate ports. If you encounter problems, see the [FortiGuard](#) discussion in the [Fortinet Communication Ports and Protocols](#) chapter of the FortiOS Handbook.

5. Is there proper routing to reach the FortiGuard servers?

Ensure there is a static or dynamic route that allows your FortiGate to reach the FortiGuard servers. Usually a generic default route to the internet is enough, but you may need to verify this if your network is complex.

6. Are there issues with DNS?

An easy way to test this is to attempt a traceroute from behind the FortiGate to an external network using the Fully Qualified Domain Name (FQDN) for a location. If the traceroute FQDN name doesn't resolve, you have general DNS problems.

7. Is there anything upstream that might be blocking FortiGuard traffic, either on the network or ISP side?

Many firewalls block all ports, by default, and ISPs often block ports that are low. There may be a firewall between the FortiGate and the FortiGuard servers that's blocking the traffic. FortiGuard uses port 53, by default, so if that port is blocked you need to either open a hole for it or change the port it is using.

8. Is there an issue with source ports?

It's possible that ports that the FortiGate uses to contact FortiGuard are being changed before they reach FortiGuard or on the return trip before they reach the FortiGate. A possible solution for this is to use a fixed-port at NAT'd firewalls to ensure the port remains the same. You can use packet sniffing to find more information about what's happening with ports.

9. Are there security policies that include antivirus?

If none of the security policies include antivirus, the antivirus database won't be updated. If antivirus is included, only the database type that's used will be updated.

FortiGuard server settings

Your local FortiGate connects to remote FortiGuard servers to get updates to FortiGuard information, such as new viruses that may have been found or other new threats. This section shows ways that you can display FortiGuard server information on your FortiGate, and how you can use that information and update it to fix potential problems.

Displaying the server list

The `get webfilter status` or `diagnose debug rating` command shows the list of FDS servers that the FortiGate uses to send web filtering requests. Rating requests are only sent to the server at the top of the list in normal operation. Each server is probed for Round Trip Time (RTT) every two minutes.

Optionally, you can add a refresh rate to the end of this command to determine how often the server list is refreshed.

Rating may not be enabled on your FortiGate.

To show the list of servers a FortiGate uses to send web filtering requests - CLI

```
get webfilter status
```

Sample output:

```
Locale : english
License : Contract
Expiration : Thu Oct 9 02:00:00 2011
== Server List (Mon Feb 18 12:55:48 2008) ==
```

IP	Weight	RTT	Flags	TZ	Packets	CurrLost	TotalLost
a.b.c.d	0	1	DI	2	1926879	0	11176
10.1.101.1	10	329		1	10263	0	633
10.2.102.2	20	169		0	16105	0	80
10.3.103.3	20	182		0	6741	0	776
10.4.104.4	20	184		0	5249	0	987
10.5.105.5	25	181		0	12072	0	178

Output details

The server list includes the IP addresses of alternate servers if the first entry can't be reached. In this example, the IP addresses are not public addresses.

The following flags in `get webfilter status` indicate the server status:

Flag	Description
D	The server was found through the DNS lookup of the hostname. If the hostname returns more than one IP address, all of them are flagged with D and are used first for INIT requests before falling back to the other servers.
I	The server to which the last INIT request was sent
F	The server hasn't responded to requests and is considered to have failed

Flag	Description
T	The server is currently being timed
S	Rating requests can be sent to the server The flag is set for a server only in two cases: <ol style="list-style-type: none">1. The server exists in the servers list received from the FortiManager or any other INIT server.2. The server list received from the FortiManager is empty so the FortiManager is the only server that the FortiGate knows and it should be used as the rating server.

Sorting the server list

The server list is sorted first by weight. The server with the smallest RTT appears at the top of the list, regardless of weight. When a packet is lost (there has been no response in 2 seconds), it's re-sent to the next server in the list. Therefore, the top position in the list is selected based on RTT, while the other positions are based on weight.

Calculating weight

The weight for each server increases with failed packets and decreases with successful packets. To lower the possibility of using a remote server, the weight isn't allowed to dip below a base weight. The base weight is calculated as the difference in hours between the FortiGate and the server multiplied by 10. The farther away the server is, the higher its base weight is and the lower it appears in the list.

Troubleshooting tips

The following tips present common causes of problems.

How to check hardware connections

- Are all of the cables and interfaces connected properly?
- Is the LED for the interface green?

How to check FortiOS network settings

- If you're having problems connecting to the management interface, is your protocol enabled on the interface for administrative access?
- Does the interface have an IP address?

How to check CPU and memory resources

- Is the CPU running at almost 100 percent usage?
- Is your FortiGate running low on memory?

How to check modem status

- Is the modem connected?
- Are there PPP issues?

How to run ping and traceroute

- Is the FortiGate experiencing complete packet loss?

How to check the logs

- Do you need to identify a problem?

How to verify the contents of the routing table (in NAT mode)

- Are there routes in the routing table for default and static routes?
- Do all connected subnets have a route in the routing table?
- Does a route have a higher priority than it should?

How to verify the correct route is being used

- Is the traffic routed correctly?

How to verify the correct firewall policy is being used

- Is the correct firewall policy applied to the expected traffic?

How to check the bridging information in transparent mode

- Are you having problems in transparent mode?

How to check number of sessions used by UTM proxy

- Have you reached the maximum number of sessions for a protocol?
- Are new sessions failing to start for a certain protocol?

How to examine the firewall session list

- Are there active firewall sessions?

[How to check wireless information](#)

- Is the wireless network working properly?

[How to verify FortiGuard connectivity](#)

- Is the FortiGate communicating properly with FortiGuard?

[How to perform a sniffer trace \(CLI and packet capture\)](#)

- Is traffic entering the FortiGate? Does the traffic arrive on the expected interface?
- Is the ARP resolution correct for the next-hop destination?
- Is the traffic exiting the FortiGate to the destination as expected?
- Is the FortiGate sending traffic back to the originator?

[How to debug the packet flow](#)

- Is traffic entering or leaving the FortiGate as expected?

How to check hardware connections

If there's no traffic flowing from the FortiGate, you may have a hardware problem.

To check hardware connections:

- Ensure network cables are plugged into the interfaces.
- Verify that the LED connection lights for the network cables are the right color (usually green).
- If the cable or its connector are damaged, change the cable. You should also change the cable if you're not sure about the type or quality of the cable, such as straight through or crossover, or if you see exposed wires at the connector.
- Connect the FortiGate to different hardware.
- Ensure the link status is set to **Up** for the interface (see **Network > Interfaces**). The link status is based on the physical connection and can't be set in FortiOS.

If any of these solve the problem, it was a hardware connection problem. You should still perform some basic software connectivity tests to ensure complete connectivity. The interface might also be disabled, or its **Administrative Status** might be set to **Down**.

To enable an interface - GUI

1. Go to **Network > Interfaces**.
2. Select and edit the interface to enable, such as **port1**.
3. Find **Administrative Status** at the bottom of the screen, and select **Up**.
4. Select **Apply**.

To enable an interface - CLI

```
config system interface
  edit port1
    set status up
  next
end
```

How to check FortiOS network settings

You can manage FortiOS network settings in the GUI and the CLI. The following information includes troubleshooting and best practice information. The network settings include:

- [Interface settings](#)
- [DNS settings](#)
- [DHCP server settings](#)

Interface settings

If you can access the FortiGate with the management cable only, the first step is to display the interface settings. To display the settings for the internal interface, use the following CLI command:

```
FGT# show system interface <interface_name>
```

For a complete listing of all the possible interface settings, use the following CLI command:

```
config system interface
  edit <interface_name>
  get
end
```

Check the interface settings to ensure that they aren't preventing traffic. Check the following items (only the GUI terms are shown, CLI terms may vary):

Setting	Description
Link Status	Down until a valid cable is plugged into this interface, then it will be Up . The Link Status is shown physically by the connection LED for the interface. If the LED is green, the connection is good. If Link Status is Down , the interface doesn't work. Link Status is also displayed on the Network > Interfaces screen, by default.
Addressing mode	Don't use DHCP if you don't have a DHCP server. You won't be able to log in to an interface in DHCP mode as it won't have an IP address.
IP/Network Mask	An interface needs an IP address to be able to connect to other devices. Ensure there is a valid IP address in this field. The one exception is if DHCP is enabled for this interface to get its IP address from an external DHCP server.
IPv6 address	The same protocol must be used by both ends to complete the connection. Ensure both this interface and the remote connection are both using IPv4 or both using IPv6 addressing.
Administrative access	If no protocols are selected, you will have to use the local management cable to connect to the unit. If you're using IPv6, configure the IPv6 administrative access protocols.
Administrative status	Set to Up or the interface won't work.

DNS settings

You can trace many networking problems back to DNS issues. Check the following items:

- Are there values for both primary and secondary entries?
- Is the local domain name correct?
- Are you using IPv6 addressing? If so, are the IPv6 DNS settings correct?
- Are you using Dynamic DNS (DDNS)? If so, is it using the correct server, credentials, and interface?
- Can you contact both DNS servers to verify the servers are operational?
- If an interface addressing mode is set to DHCP and is set to override the internal DNS, is that interface receiving a valid DNS entry from the DHCP server? Is it a reasonable address and can it be contacted to verify it's operational?
- Are there any DENY security policies that need to allow DNS?
- Can any internal device perform a successful traceroute to a location using the FQDN?

DHCP server settings

DHCP servers are common on internal and wireless networks. If the DHCP server isn't configured correctly, it can cause problems. Check the following items:

- Is the DHCP server entry set to **Relay**? If so, verify there is another DHCP server to which requests can be relayed. Otherwise, it should be set to **Server**.
- Is the DHCP server enabled?
- Does the DHCP server use a valid IP address range? Are other devices using the addresses? If one or more devices are using IP addresses in this range, you can use the IP reservation feature to ensure the DHCP server doesn't use these addresses.
- Is there a gateway entry? If not, add a gateway entry to ensure that the server's clients have a default route.
- Is the system DNS setting being used? A best practice is to avoid confusion by using the system DNS whenever possible. However, you can specify up to three custom DNS servers, and you should use all three entries for redundancy.



There are some situations, such as a new wireless interface, or during the initial FortiGate configuration, where interfaces override the system DNS entries. When this happens, it often shows up as intermittent Internet connectivity. To fix the problem, go to **Network > DNS** and enable **Use FortiGuard Servers**.

How to check CPU and memory resources

System resources are shared and a number of processes run simultaneously on the FortiGate unit.

The **Resource** widgets for **CPU** and **Memory** on the **Dashboard** provide a quick way to monitor usage.

To use the CLI to check the system resources on your FortiGate unit, run the following command:

```
FGT# get system performance status
```

This command provides a quick and easy snapshot of the FortiGate.

The first line of output shows the CPU usage by category. A FortiGate that is doing nothing looks like the following example:

```
CPU states: 0% user 0% system 0% nice 100% idle
```

However, if your network is running slow you might see something looks like the following example:

```
CPU states: 1% user 98% system 0% nice 1% idle
```

This line shows that all the CPU is used up by system processes. Normally this shouldn't happen since it shows that the FortiGate is overloaded for some reason. If you see this overloading, you should investigate further because it's possible a process, such as `scanunitid`, is using all the resources to scan traffic, in which case you need to reduce the amount of traffic that's being scanned by blocking unwanted protocols, configuring more security policies to limit scanning to certain protocols, or similar actions. It's also possible that a hacker has gained access to your network and is overloading it with malicious activity, such as running a spam server or using zombie PCs to attack other networks on the Internet. You can use the `get system performance top` CLI command to get more information about the CPU. This command shows all of the top processes that are running on the FortiGate (the names are on the left) and their CPU usage. If a process is using most of the CPU cycles, investigate it to determine whether the activity is normal.

The second line of output from the `get system performance status` command shows the memory usage. Memory usage shouldn't exceed 90%. If memory is too full, some processes won't be able to function properly. For example, if the system is running low on memory, antivirus scanning enters into failopen mode where it drops connections or bypasses the antivirus system.

The other lines of output, such as average network usage, average session setup rate, viruses caught, and IPS attacks blocked, can also help you determine why system resource usage is high. For example, if network usage is high it results in high traffic processing on the FortiGate, or if the session setup rate is very low (or zero) the proxy may be overloaded and unable to do its job.

How to troubleshoot high memory usage

As with any system, a FortiGate has limited hardware resources, such as memory, and all processes running on the FortiGate share the memory. Each process uses more or less memory, depending on its workload. For example, a process usually uses more memory in high traffic situations. If some processes use all of the available memory, other processes won't be able to run.

When high memory usage occurs, the services may freeze up, connections may be lost, or new connections may be refused.

If you see high memory usage in the **Memory** widget, the FortiGate may be handling high traffic volumes. Alternatively, the FortiGate may have with connection pool limits that are affecting a single proxy. If the FortiGate receives large volumes of traffic on a specific proxy, the unit may exceed the connection pool limit. If the number of free connections within a proxy connection pool reaches zero, issues may occur.

Use the following CLI command, which uses the antivirus failopen feature:

```
config system global
    set av-failopen idledrop
end
```

If you set `av-failopen` to `idledrop`, the FortiGate drops connections based on the clients that have the most open connections. This helps you determine the behavior of the FortiGate antivirus system if it becomes overloaded in high traffic.

Use the following CLI command, which shows information about current memory usage:

```
diagnose hardware sysinfo memory
```

Sample output:

```

total: used: free: shared: buffers: cached: shm:
Mem: 2074185728 756936704 1317249024 0 20701184 194555904 161046528
Swap:      0      0      0
MemTotal:   2025572 kB
MemFree:    1286376 kB
MemShared:      0 kB
Buffers:     20216 kB
Cached:      189996 kB
SwapCached:    0 kB
Active:       56644 kB
Inactive:    153648 kB
HighTotal:      0 kB
HighFree:      0 kB
LowTotal:     2025572 kB
LowFree:      1286376 kB
SwapTotal:      0 kB
SwapFree:      0 kB

```

How to troubleshoot high CPU usage

If you deploy too many FortiOS features at the same time, you can overextend the CPU resources on a FortiGate. When this occurs, the FortiGate undergoes connection-related problems.

Some examples of CPU intensive features are:

- VPN high-level encryption
- Intensive scanning of all traffic
- Logging all traffic and packets
- Dashboard widgets that frequently perform data updates

1. Determine the current level of CPU usage.

There are two ways to do this. The simplest is to look at the **CPU** widget on the **Dashboard** in the FortiGate GUI. The real-time CPU usage is displayed for different timeframes. The second method provides precise usage values both for overall usage and for specific processes. To use it, run the `diagnose system top` CLI command.

Sample output:

```

Run Time: 86 days, 0 hours and 10 minutes
OU, ON, OS, 100I, OWA, OHI, OSI, OST; 3040T, 2437F
bcm.user 93 S < 3.1 0.4
httpsd 18922 S 1.5 0.5
httpsd 19150 S 0.3 0.5
newcli 20195 R 0.1 0.1
cmdbsvr 115 S 0.0 0.8
pyfcgid 20107 S 0.0 0.6
forticron 146 S 0.0 0.5
httpsd 139 S 0.0 0.5
cw_acd 166 S 0.0 0.5
miglogd 136 S 0.0 0.5
pyfcgid 20110 S 0.0 0.4
pyfcgid 20111 S 0.0 0.4
pyfcgid 20109 S 0.0 0.4

```



```

httpsd 20192 S 0.0 0.4
miglogd 174 S 0.0 0.4
miglogd 175 S 0.0 0.4
fgfmd 165 S 0.0 0.3
newcli 20191 S 0.0 0.3
initXXXXXXXXXX 1 S 0.0 0.3
httpsd 184 s 0.0 0.3

```

The codes displayed on the second output line are explained in the following table:

Code	Description
U	the percentage of user space applications that are currently using the CPU
N	the percentage of time that the CPU spent on low priority processes since the last shutdown
S	the percentage of system processes (or kernel processes) that are using the CPU
I	the percentage of idle CPU resources
WA	the percentage of time that the CPU spent waiting on IO peripherals since the last shutdown
HI	the percentage of time that the CPU spent handling hardware interrupt routines since the last shutdown
SI	the percentage of time that the CPU spent handling software interrupt routines since the last shutdown
ST	or steal time, is the percentage of time that a virtual CPU waits for the physical CPU when the hypervisor is servicing another virtual processor
T	is the total FortiOS system memory, in MB
F	is free memory, in MB

Each additional line of the command output displays information for specific processes running on the FortiGate unit. For example, the third line of the output is:

```
newcli 20195 R 0.1 0.1
```

The items in the third line of output are explained in the following table:

Item	Description
newcli	the process name Other process names can include <code>ipsengine</code> , <code>sshd</code> , <code>cmdbsrv</code> , <code>httpsd</code> , <code>scanunitd</code> , and <code>miglogd</code> . Duplicate process names indicate that separate instances of that process that are running.
20195	the process ID, which can be any number
R	current state of the process. The process state can be: <ul style="list-style-type: none"> • R - running • S - sleep • Z - zombie • D - disk sleep
0.1	the percentage of CPU capacity that the process is using CPU usage can range from 0.0 for a process that is sleeping to higher values for a process that's taking a lot of CPU time
0.1	the amount of memory that the process is using Memory usage can range from 0.1 to 5.5 and higher.

When `diagnose system top` is running, enter the following single-key commands:

- `q` to quit and return to the normal CLI prompt.
- `p` to sort the processes by the amount of CPU that the processes are using.
- `m` to sort the processes by the amount of memory that the processes are using.

The processes listed are the top processes that are running not all of the processes. For example, if there are 20 processes listed, they are the top 20 processes, sorted by either CPU or memory usage. You can configure the number of processes displayed, using the following CLI command:

```
diagnose system top <integer_seconds> <integer_maximum_lines>
```

Where:

- `<integer_seconds>` is the delay in seconds (default is 5)
 - `<integer_maximum_lines>` is the maximum number of lines (or processes) to list (default is 20)
2. Determine what features are using most of the CPU resources.
There is a command in the CLI that shows the top few processes that are currently running and use the most CPU resources. The `get system performance top` command shows a table of information. The second column from the right shows CPU usage by percentage. If the top few entries are using most of the CPU, note which processes they are and try to reduce their CPU load.

Some examples of processes you'll see include:

- `ipsengine` — the IPS engine that scans traffic for intrusions
- `scanunitd` — antivirus scanner
- `httpsd` — secure HTTP
- `iked` — internet key exchange (IKE) in use with IPsec VPN tunnels
- `newcli` — active whenever you're accessing the CLI
- `sshd` — there are active secure socket connections
- `cmdbsrv` — the command database server application

Go to the features that are at the top of the list and look for evidence of them overusing the CPU. Generally, the monitor for a feature is a good place to start.

3. Check for unnecessary CPU “wasters”.

These are some best practices that will reduce your CPU usage, even if the FortiGate isn't experiencing high CPU usage. Note that if the following information instructs you to turn off a feature that you require, disregard that part of the instructions.

- Use hardware acceleration wherever possible to offload tasks from the CPU. Offloading tasks, such as encryption, frees up the CPU for other tasks.
- Avoid the use of GUI widgets that require computing cycles, such as the Top Sessions widget. These widgets constantly poll the system for information, which uses CPU and other resources.
- Schedule antivirus, IPS, and firmware updates during off peak hours. These updates don't usually consume CPU resources but they can disrupt normal operation.
- Check the log levels and which events are being logged. This is the severity of the messages that are recorded. Consider going up one level to reduce the amount of logging. Also, if there are events you don't need to monitor, remove them from the list.
- Log to FortiCloud instead of logging to memory or disk. Logging to memory quickly uses up resources and logging to local disk impacts overall performance and reduces the lifetime of the unit. Fortinet recommends that you log to FortiCloud because it doesn't use much CPU.
- If the disk is almost full, transfer the logs or data off the disk to free up space. When a disk is almost full it consumes a lot of resources to find free space and organize the files.
- If packet logging is enabled on the FortiGate, consider disabling it. When packet logging is enable, it records every packet that comes through that policy.
- Halt all sniffers and traces.
- Ensure the FortiGate isn't scanning traffic twice. If traffic enters the FortiGate on one interface, goes out another, and then comes back in again, that traffic doesn't need to be rescanned. Doing so is a waste of resources. However, ensure that traffic truly is being scanned once.
- Reduce the session timers to close unused sessions faster. Enter the following CLI commands, which reduce the default values. Note that, by default, the system adds 10 seconds to `tcp-timewait`.

```
config system global
    set tcp-halfclose-timer 30
    set tcp-halfopen-timer 30
    set tcp-timewait-timer 0
    set udp-idle-timer 60
end
```

- In **System > Feature Visibility**, only enable features that you need.

4. When CPU usage is under control, use SNMP to monitor CPU usage. Alternatively, use logging to record CPU and memory usage every 5 minutes.

Once things are back to normal, you should set up a warning system that sends you alerts when CPU resources

are used excessively. A common method to do this is using SNMP. SNMP monitors many values in FortiOS and allows you to set high water marks that generate events. You run an application on your computer to watch for and record these events. Go to **System > SNMP** to enable and configure an SNMP community. If this method is too complicated, you can use the **System Resources** widget to record CPU usage. However, this method only records problems as they happen and won't send you alerts for problems.

How to check modem status

If the modem doesn't work properly or a FortiGate doesn't detect the modem, you can use the following diagnostic command to help you troubleshoot issues with the modem:

```
diagnose system modem {cmd | com | detect | history | external-modem | query| reset}
```

You should always run the following diagnose command after you insert the USB modem into the FortiGate:

```
diagnose system modem detect
```

You can view the modem configuration, using the `get system modem` command. You can also view the modem's vendor and the custom product identification number from the information output from the `get system modem` command.

When there are connectivity issues, use the following to help you resolve them:

- `diagnose debug enable` – activates the debug on the console
- `diagnose debug application modemd` – dumps communication between the modem and the unit.
- `diagnose debug application ppp` – dumps the PPP negotiating messages.
- `execute modem dial` – displays modem debug output.

The modem diagnose output shouldn't contain errors on the way to initializing. You should also verify the number that is used to dial into your ISP.

How to run ping and traceroute

Ping and traceroute are useful tools in network troubleshooting. Alone, either one can determine network connectivity between two points. However, ping can be used to generate simple network traffic that you can view using `diagnose` commands on the FortiGate. This combination can be very powerful when you're trying to locate network problems.

In addition to their normal uses, ping and traceroute can tell you if your computer or network device has access to a domain name server (DNS). While both tools can use IP addresses alone, they can also use domain names for devices. This is an added troubleshooting feature that can be useful in determining why particular services, such as email or web browsing, may not work properly.



If ping doesn't work, it may be disabled on at least one of the interface settings, and security policies for that interface.

Both ping and traceroute require particular ports to be open on firewalls, or else they can't function. Since you typically use these tools to troubleshoot, you can allow them in the security policies and on interfaces only when you need them. Otherwise, keep the ports disabled for added security.

Ping

The ping command sends a very small packet to a destination, and waits for a response. The response has a timer that expires when the destination is unreachable.

Ping is part of layer 3 on the OSI Networking Model. Ping sends Internet Control Message Protocol (ICMP) “echo request” packets to the destination, and listens for “echo response” packets in reply. However, many public networks block ICMP packets because ping can be used in a denial of service (DoS) attack (such as Ping of Death or a smurf attack), or by an attacker to find active locations on the network. By default, FortiGate units have ping enabled while broadcast-forward is disabled on the external interface.

What ping can tell you

Beyond the basic connectivity information, ping can tell you the amount of packet loss (if any), how long it takes the packet to make the round trip, and the variation in that time from packet to packet.

If there is some packet loss detected, you should investigate the following:

- Possible ECMP, split horizon, or network loops.
- Cabling, to ensure no loose connections.
- Verify which security policy was used (use the packet count column on the **Policy & Objects > IPv4 Policy** or **Policy & Objects > IPv6 Policy** page).

If there is total packet loss, you should investigate the following:

- **Hardware:** ensure cabling is correct, and all equipment between the two locations is accounted for.
- **Addresses and routes:** ensure all IP addresses and routing information along the route is configured as expected.
- **Firewalls:** ensure all firewalls, including FortiGate security policies allow PING to pass through.

How to use ping

Ping syntax is the same for nearly every type of system on a network.

To ping from a FortiGate unit

1. Connect to the CLI either through telnet or through the CLI widget on the **Dashboard**.
2. Enter `exec ping 10.11.101.101` to send 5 ping packets to the destination IP address. There are no options for this command.

Sample output:

```
Head_Office_620b # exec ping 10.11.101.101
PING 10.11.101.101 (10.11.101.101): 56 data bytes
64 bytes from 10.11.101.101: icmp_seq=0 ttl=255 time=0.3 ms
64 bytes from 10.11.101.101: icmp_seq=1 ttl=255 time=0.2 ms
64 bytes from 10.11.101.101: icmp_seq=2 ttl=255 time=0.2 ms
64 bytes from 10.11.101.101: icmp_seq=3 ttl=255 time=0.2 ms
64 bytes from 10.11.101.101: icmp_seq=4 ttl=255 time=0.2 ms

--- 10.11.101.101 ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 0.2/0.2/0.3 ms
```

To ping from a MicroSoft Windows PC

1. Open a command window.
2. Enter `ping 10.11.101.100` to ping the default internal interface of the FortiGate with four packets.

Other options include:

- `-t` to send packets until you press "Ctrl+C"
- `-a` to resolve addresses to domain names where possible
- `-n X` to send X ping packets and stop

Sample output:

```
C:\>ping 10.11.101.101

Pinging 10.11.101.101 with 32 bytes of data:
Reply from 10.11.101.101: bytes=32 time=10ms TTL=255
Reply from 10.11.101.101: bytes=32 time<1ms TTL=255
Reply from 10.11.101.101: bytes=32 time=1ms TTL=255
Reply from 10.11.101.101: bytes=32 time=1ms TTL=255

Ping statistics for 10.11.101.101:
    Packets: Sent = 4, Received = 4, Lost = 0 (0% loss),
    Approximate round trip times in milli-seconds:
        Minimum = 0ms, Maximum = 10ms, Average = 3ms
```

To ping from a Linux PC

1. Go to a shell prompt.
2. Enter `"ping 10.11.101.101"`.

Traceroute

Where ping will only tell you if it reached its destination and returned successfully, traceroute shows each step of the journey to its destination and how long each step takes. If ping finds an outage between two points, you can use traceroute to locate exactly where the problem is.

What is traceroute

Traceroute works by sending ICMP packets to test each hop along the route. It sends three packets, and then increases the time to live (TTL) setting by one each time. This effectively allows the packets to go one hop farther along the route. This is the reason why most traceroute commands display their maximum hop count before they start tracing the route, which is the maximum number of steps it takes before it declares the destination unreachable. Also, the TTL setting may result in steps along the route timing out due to slow responses. There are many possible reasons for this to occur.

By default, traceroute uses UDP datagrams with destination ports numbered from 33434 to 33534. The traceroute utility may also offer the option to select use of ICMP echo request (type 8) instead, which the Windows `tracert` utility uses. If you must, allow both protocols inbound through the FortiGate security policies (UDP with ports from 33434 to 33534 and ICMP type 8).

You can also use the packet count column of the **Policy & Objects > IPv4 Policy** (or, **Policy & Objects > IPv6 Policy**, if applicable) page to track traceroute packets. This allows you to verify the connection and confirm which security policy the traceroute packets are using.

What traceroute can tell you

Ping and traceroute have similar functions, which is to verify connectivity between two points. The big difference is that traceroute shows you each step of the way, where ping doesn't. Also, ping and traceroute use different protocols and ports, so one may succeed where the other fails.

You can verify your DNS connection using traceroute. If you enter an FQDN instead of an IP address for the traceroute, DNS tries to resolve that domain name. If the name isn't resolved, you have DNS issues.

How to use traceroute

The traceroute command varies slightly between operating systems. Note that in MicroSoft Windows, the command name is shortened to "tracert". Also, your output lists different domain names and IP addresses along your route.

To use traceroute on a MicroSoft Windows PC

1. Open a command window.
2. Enter "tracert fortinet.com" to trace the route from the PC to the Fortinet web site.

Sample output:

```
C:\>tracert fortinet.com

Tracing route to fortinet.com [208.70.202.225]
over a maximum of 30 hops:
 1 <1 ms <1 ms <1 ms 172.20.120.2
 2 66 ms 24 ms 31 ms 209-87-254-xxx.storm.ca [209.87.254.221]
 3 52 ms 22 ms 18 ms core-2-g0-0-1104.storm.ca [209.87.239.129]
 4 43 ms 36 ms 27 ms core-3-g0-0-1185.storm.ca [209.87.239.222]
 5 46 ms 21 ms 16 ms te3-x.1156.mpd01.cogentco.com [38.104.158.69]
 6 25 ms 45 ms 53 ms te8-7.mpd01.cogentco.com [154.54.27.249]
 7 89 ms 70 ms 36 ms te3-x.mpd01.cogentco.com [154.54.6.206]
 8 55 ms 77 ms 58 ms sl-st30-chi-.sprintlink.net [144.232.9.69]
 9 53 ms 58 ms 46 ms sl-0-3-3-x.sprintlink.net [144.232.19.181]
10 82 ms 90 ms 75 ms sl-x-12-0-1.sprintlink.net [144.232.20.61]
11 122 ms 123 ms 132 ms sl-0-x-0-3.sprintlink.net [144.232.18.150]
12 129 ms 119 ms 139 ms 144.232.20.7
13 172 ms 164 ms 243 ms sl-321313-0.sprintlink.net [144.223.243.58]
14 99 ms 94 ms 93 ms 203.78.181.18
15 108 ms 102 ms 89 ms 203.78.176.2
16 98 ms 95 ms 97 ms 208.70.202.225

Trace complete.
```

The first column on the left is the hop count, which can't exceed 30 hops. When that number is reached, the traceroute ends.

The second, third, and fourth columns display how much time each of the three packets takes to reach this stage of the route. These values are in milliseconds and normally vary quite a bit. Typically a value of <1ms indicates a local connection.

The fifth column (farthest to the right) shows the domain name of the device and its IP address, or possibly only the IP address.

To perform a traceroute on a Linux PC

1. Go to a command line prompt.
2. Enter "traceroute fortinet.com".

The Linux traceroute output is very similar to the MicroSoft Windows tracert output.

To trace a route from a FortiGate to a destination IP address - CLI

```
# execute traceroute www.fortinet.com
traceroute to www.fortinet.com (66.171.121.34), 32 hops max, 84 byte packets
 1 172.20.120.2 0.637 ms 0.653 ms 0.279 ms
 2 209.87.254.221 <static-209-87-254-221.storm.ca> 2.448 ms 2.519 ms 2.458 ms
 3 209.87.239.129 <core-2-g0-2.storm.ca> 2.917 ms 2.828 ms 9.324 ms
 4 209.87.239.199 <core-3-bdi1739.storm.ca> 13.248 ms 12.401 ms 13.009 ms
 5 216.66.41.113 <v502.core1.tor1.he.net> 17.181 ms 12.422 ms 12.268 ms
 6 184.105.80.9 <100ge1-2.core1.nyc4.he.net> 21.355 ms 21.518 ms 21.597 ms
 7 198.32.118.41 <ny-paix-gni.twgate.net> 83.297 ms 84.416 ms 83.782 ms
 8 203.160.228.217 <217-228-160-203.TWGATE-IP.twgate.net> 82.579 ms 82.187 ms 82.066 ms
 9 203.160.228.229 <229-228-160-203.TWGATE-IP.twgate.net> 82.055 ms 82.455 ms 81.808 ms
10 203.78.181.2 82.262 ms 81.572 ms 82.015 ms
11 203.78.186.70 83.283 ms 83.243 ms 83.293 ms
12 66.171.127.177 84.030 ms 84.229 ms 83.550 ms
13 66.171.121.34 <www.fortinet.com> 84.023 ms 83.903 ms 84.032 ms
14 66.171.121.34 <www.fortinet.com> 83.874 ms 84.084 ms 83.810 ms
```

How to check the logs

You might forget this step in troubleshooting, but it's an important one. Logging records the traffic that passes through the FortiGate to your network and what action the FortiGate takes when it scans the traffic. This information record is called a log message.

When you first configure FortiOS, log as much information as you can. If the logs that the FortiGate generates are too large, you can turn off or scale back the logging for features that you're not using.

As with most troubleshooting steps, before you can determine if the logs indicate a problem, you need to know what logs result from normal operation. Without a baseline it's difficult to troubleshoot.

When you troubleshoot with log files:

- Compare current logs to a recorded baseline of normal operation.
- If you need to, increase the level of logging (such as from Warning to Information) to obtain more information.

When increasing logging levels, ensure that you configure email alerts and select both disk usage and log quota. This ensures that you'll be notified if the increase in logging causes problems. Configure log settings by going to **Log & Report > Log Settings**.

Determine the activities that generate the most log entries:

- Check all logs to ensure important information isn't overlooked.
- Filter or order log entries based on different fields, such as level, service, or IP address, to look for patterns that may indicate a specific problem, such as frequent blocked connections on a specific port for all IP addresses.

Logs will help identify and locate any problems, but they will not solve the problems. The purpose of logs is to speed up your problem solving and save you time and effort.

For more information about logging and log reports, see the [Logging and Reporting handbook chapter](#).

How to verify the contents of the routing table (in NAT mode)

When a FortiGate has limited or no connectivity, a good place to look for information is the routing table.

The routing table is where all the currently used routes are stored for both static and dynamic protocols. If a route is in the routing table, it saves time and resources that you would spend performing a lookup. If a route is not used for a while and a new route needs to be added, the oldest least used route is bumped if the routing table is full. This ensures that the most recently used routes stay in the table. If your FortiGate is in transparent mode, you can't perform this step.

If the FortiGate is running in NAT mode, verify that all desired routes are in the routing table, including local subnets, default routes, specific static routes, and dynamic routing protocols.

To check the routing table in the FortiGate GUI, use the Routing Monitor by going to **Monitor > Routing Monitor**.

In the CLI, use the command `get router info routing-table all`.

Sample output:

```
FGT# get router info routing-table all

Codes:
  K - kernel, C - connected, S - static, R - RIP, B - BGP
  O - OSPF, IA - OSPF inter area
  N1 - OSPF NSSA external type 1, N2 - OSPF NSSA external type 2
  E1 - OSPF external type 1, E2 - OSPF external type 2
  i - IS-IS, L1 - IS-IS level-1, L2 - IS-IS level-2, ia - IS-IS inter area
  * - candidate default

S* 0.0.0.0/0 [10/0] via 172.20.120.2, wan1
C 10.31.101.0/24 is directly connected, internal
C 172.20.120.0/24 is directly connected, wan1
```

How to verify the correct route is being used

If you have more than one default route and want to make sure that traffic is flowing as expected and through the right route, you can run a trace route from a machine in the local area network (LAN). This shows you the first hop that the traffic goes through.

Sample output:

```
C:\>tracert www.fortinet.com

Tracing route to www.fortinet.com [66.171.121.34]
over a maximum of 30 hops:

 1 <1 ms <1 ms <1 ms 10.10.1.99
 2 1 ms <1 ms <1 ms 172.20.120.2
 3 3 ms 3 ms 3 ms static-209-87-254-221.storm.ca [209.87.254.221]
 4 3 ms 3 ms 3 ms core-2-g0-2.storm.ca [209.87.239.129]
 5 13 ms 13 ms 13 ms core-3-bdi1739.storm.ca [209.87.239.199]
 6 12 ms 19 ms 11 ms v502.core1.tor1.he.net [216.66.41.113]
 7 22 ms 22 ms 21 ms 100ge1-2.core1.nyc4.he.net [184.105.80.9]
 8 84 ms 84 ms 84 ms ny-paix-gni.twgate.net [198.32.118.41]
 9 82 ms 84 ms 82 ms 217-228-160-203.TWGATE-IP.twgate.net [203.160.22
 8.217]
10 82 ms 81 ms 82 ms 229-228-160-203.TWGATE-IP.twgate.net [203.160.22
 8.229]
11 82 ms 82 ms 82 ms 203.78.181.2
12 84 ms 83 ms 83 ms 203.78.186.70
13 84 ms * 85 ms 66.171.127.177
14 84 ms 84 ms 84 ms fortinet.com [66.171.121.34]
15 84 ms 84 ms 83 ms fortinet.com [66.171.121.34]

Trace complete.
```

In this scenario, the first hop contains the IP address `10.10.1.99`, which is the internal interface of the FortiGate. The second hop contains the IP address `172.20.120.2`, to which the `wan1` interface of the FortiGate is connected, so we can conclude that the route through `wan1` interface is being used for this traffic.

You can also see the route taken for each session by debugging the packet flow in the CLI. For more information, see [How to debug the packet flow](#).

How to verify the correct firewall policy is being used

If you have more than one firewall policy, use the count column to check which policy is being used, the count must show traffic increasing. To do so, go to the **Policy & Objects > IPv4 Policy** or **Policy & Objects > IPv6 Policy** page.

Also, debugging the packet flow in the CLI shows the policy ID that's allowing the traffic. For more information about debugging the packet flow, see [How to debug the packet flow](#).

How to check the bridging information in transparent mode

When the FortiGate is set to transparent mode, it acts like a bridge and sends all incoming traffic out on the other interfaces. The bridge is between interfaces on the FortiGate.

Each bridge that's listed is a link between interfaces. Where traffic is flowing between interfaces, you can expect to find bridges listed. If you're having connectivity issues, and there are no bridges listed, that is a likely cause. Check for the MAC address of the interface or device in question.

How to check the bridging information

To list the existing bridge instances on the FortiGate, use the following command:

```
diagnose netlink brctl list
```

Sample output:

```
#diagnose netlink brctl list
list bridge information
1. root.b fdb: size=256 used=6 num=7 depth=2 simple=no
Total 1 bridges
```

How to display forwarding domain information

You can use forwarding domains, or collision domains, in routing to limit where packets are forwarded on the network. Layer 2 broadcasts are limited to the same group. By default, all interfaces are in group 0. For example, if the FortiGate has 12 interfaces, only two may be in the same forwarding domain, which limits packets that are broadcast to those two interfaces. This reduces traffic on the rest of the network.

Collision domains prevent the forwarding of ARP packets to all VLANs on an interface. Without collision domains, duplicate MAC addresses on VLANs may cause ARP packets to be duplicated. Duplicate ARP packets can cause some switches to reset. It's important to know what interfaces are part of which forwarding domains because this determines which interfaces can communicate with each other.

To manually configure forwarding domains in transparent mode, use the following CLI command:

```
config system interface
  edit <interface_name>
    set forward-domain <integer>
  end
```

To display the information for forward domains, use the following command:

```
diagnose netlink brctl domain <name> <id>
```

where <name> is the name of the forwarding domain to display and <id> is the domain ID.

Sample output

```
diagnose netlink brctl domain ione 101
show bridge root.b ione forward domain.
id=101 dev=trunk_1 6
```

To list the existing bridge MAC table, use the following command:

```
diagnose netlink brctl name host <name>
```

Sample output

```
show bridge control interface root.b host.
fdb: size=256, used=6, num=7, depth=2, simple=no
Bridge root.b host table
```

port no	device	devname	mac addr	ttl	attributes
2	7	wan2	02:09:0f:78:69:00	0	Local Static
5	6	vlan_1	02:09:0f:78:69:01	0	Local Static
3	8	dmz	02:09:0f:78:69:01	0	Local Static
4	9	internal	02:09:0f:78:69:02	0	Local Static
3	8	dmz	00:80:c8:39:87:5a	194	
4	9	internal	02:09:0f:78:67:68	8	
1	3	wan1	00:09:0f:78:69:fe	0	Local Static

To list the existing bridge port list, use the following command:

```
diagnose netlink brctl name port <name>
```

Sample output:

```
show bridge root.b data port.
trunk_1 peer_dev=0
internal peer_dev=0
dmz peer_dev=0
wan2 peer_dev=0
wan1 peer_dev=0
```

How to check the number of sessions that UTM proxy uses

Each FortiGate model has a limit for the maximum number of sessions that the UTM proxy supports. The UTM proxy handles all the traffic for the following protocols: HTTP, SMTP, POP3, IMAP, FTP, and NNTP. If the proxy

for a protocol fills up its session table, the FortiGate enters conserve mode, where it behaves differently, until entries and memory free up again.

Conserve or failopen mode

Once you reach the limit, depending on the conserve mode configuration on the FortiGate, no new sessions are created until old ones end. You can configure your the behavior FortiGate when memory is running low or the proxy connection limit has been reached. There are two related commands for this in the CLI:

```
config system global
  set av-failopen-session {enable | disable}
  set av-failopen { idledrop | off | one-shot | pass}
end
```

To set the behavior for these conditions, you must enable `av-failopen-session`. When it's enabled, and a proxy for a protocol runs out of room in its session table, that protocol goes into failopen mode and behaves as defined in the `av-failopen` command.

`av-failopen` determines the behavior of the proxy until entries are free in the session table again for that proxy.

- **idledrop** — This option removes idle sessions from the session table, starting with the clients that have the most sessions currently open. This method assumes that idle sessions aren't being used and it won't cause problems to close these sessions. This is usually true, but some applications may have problems opening a session. If this occurs, try another method to check if this is really the problem. This is a secure option as no unscanned traffic is allowed to pass.
- **off** — This option turns off accepting any new AV sessions, but continues to process any existing AV sessions that are currently active. All the protocols listed (HTTP, SMTP, POP3, IMAP, FTP, and NNTP) are scanned by FortiGate Antivirus. If AV scanning is enabled, `av-failopen off` is set, and the proxy session table fills up, which means that no new sessions of that type are accepted. For example, if the POP3 session table is full and email AV scanning is enabled, no additional POP3 connections are allowed until the session table gets some free space. This is a secure option because no unscanned traffic is allowed to pass.
- **one-shot** — When memory is low, bypass the antivirus system. The term `one-shot` comes from the fact that once you're in `one-shot av-failopen` mode, you must set `av-failopen` to either `pass` or `off` to restart AV scanning. This is a very unsecure option because it allows all traffic without AV scanning, and it never reverts to normal without manual assistance.
- **pass** — When memory is low, bypass the antivirus system. The difference between `pass` and `one-shot` options is that when memory is freed up, the system automatically starts AV scanning again. This is an unsecure option because it allows traffic to pass without AV scanning. However, it's better than `one-shot` because it automatically restarts AV scanning, when possible.

If the proxy session table is full for one or more protocols and your FortiGate enters into conserve or failopen mode, it appears as though the FortiGate has lost connection, network services are intermittent or don't exist, and yet other services work normally for a while until their sessions end and they join the queue of session-starved applications.

Checking sessions in use

To make troubleshooting this type of problem easier, sessions are broken down by which protocol they use. This provides you with statistics and errors specific to one of the protocols.



Due to the amount of output from this command, you should connect to the CLI with a terminal program, such as puTTY, that logs output. Otherwise, you may not be able to access all the output information from the command.

In the following output, only the HTTP entries are displayed. The other protocols have been removed in an attempt to shorten the output. There will be separate entries for each supported protocol (HTTP, SMTP, POP3, IMAP, FTP, and NNTP) in each section of the output.

To check sessions in use and related errors - CLI

```
FGT# # get test proxyworker 4

Worker[0]
HTTP Common
Current Connections 8/8032
Max Concurrent Connections 76

Worker Stat
Running time (HH:MM:SS:usec) 29:06:27:369365
Time in loop scanning 2:08:000198
Error Count (accept) 0
Error Count (read) 0
Error Count (write) 0
Error Count (poll) 0
Error Count (alloc) 0
Last Error 0
Acceptor Read 6386
Acceptor Write 19621
Acceptor Close 0

HTTP Stat
Bytes sent 667012 (kb)
Bytes received 680347 (kb)
Error Count (alloc) 0
Error Count (accept) 0
Error Count (bind) 0
Error Count (connect) 0
Error Count (socket) 0
Error Count (read) 134
Error Count (write) 0
Error Count (retry) 40
Error Count (poll) 0
Error Count (scan reset) 2
Error Count (urlfilter wait) 3
Last Error 104
Web responses clean 17950
Web responses scan errors 23
Web responses detected 16
Web responses infected with worms 0
Web responses infected with viruses 0
Web responses infected with susp 0
Web responses file blocked 0
Web responses file exempt 0
Web responses bannedword detected 0
```

```
Web requests oversize pass 16
Web requests oversize block 0
Last Server Scan errors 102
URL requests exempt 0
URL requests blocked 0
URL requests passed 0
URL requests submit error 0
URL requests rating error 0
URL requests rating block 0
URL requests rating allow 10025
URL requests infected with worms 0
Web requests detected 0
Web requests file blocked 0
Web requests file exempt 0
POST requests clean 512
POST requests scan errors 0
POST requests infected with viruses 0
POST requests infected with susp 0
POST requests file blocked 0
POST requests bannedword detected 0
POST requests oversize pass 0
POST requests oversize block 0
Web request backlog drop 0
Web response backlog drop 0

Worker Accounting
poll=721392/649809/42 pollfail=0 cmdb=85 scan=19266 acceptor=25975

HTTP Accounting
setup_ok=8316 setup_fail=0 conn_ok=0 conn_inp=8316
urlfilter=16553/21491/20 uf_lookupf=0
scan=23786 clt=278876 srv=368557

SMTP Accounting
setup_ok=12 setup_fail=0 conn_ok=0 conn_inp=12
scan=12 suspend=0 resume=0 reject=0 spamadd=0 spamdel=0 clt=275 srv=279

POP3 Accounting
setup_ok=30 setup_fail=0 conn_ok=0 conn_inp=30
scan=3 clt=5690 srv=5836

IMAP Accounting
setup_ok=0 setup_fail=0 conn_ok=0 conn_inp=0
scan=0 clt=0 srv=0

FTP Accounting
setup_ok=0 setup_fail=0 conn_ok=0 conn_inp=0
scan=0 clt=0 srv=0 datalisten=0 dataclt=0 datasrv=0

NNTP Accounting
setup_ok=0 setup_fail=0 conn_ok=0 conn_inp=0
scan=0 clt=0 srv=0
```

The output from this command falls into the following sections:

- **HTTP Common current connections** — This displays an entry for each protocol that displays the connections currently used, and the maximum connections allowed. This maximum is for the UTM proxy, which means all of the protocols connections combined can't be larger than this number. To support this, note that the maximum session count for each protocol is the same. You may also see a line titled `Max Concurrent Connections` for each protocol. This number is the maximum connections of this type that's allowed at one time. If VDOMs are enabled, this value is defined either on the global or per-VDOM level at **VDOM > Global Resources**.
- **Worker Stat** — This displays statistics about the UTM proxy including how long it has been running, and how many errors it has found.
- **HTTP Stat** — This section includes statistics about the HTTP protocol proxy. This is a very extensive list that includes errors, web responses, and any UTM positive matches. There are similar sections for each protocol, but the specific entries in each vary based on what UTM scanning is looking for in each — spam control for email, file transfer blocking for FTP, and so on.
- **Worker Accounting** — Lists accounting information about the UTM proxy such as polling statistics, how many sessions were scanned, and how many were just accepted. This information can show you if expected AV scanning is taking place or not. Under normal operation there shouldn't be errors or fails.
- **HTTP Accounting** — The accounting sections for each protocol provide information about successful session creation, failures, how many sessions are being scanned or filtered, and how many are client or server originated. If `setup_fail` is larger than zero, run the command again to see if it's increasing quickly. If it is, your FortiGate may be in conserve mode.

Related commands

To dump memory usage:

```
# get test proxyworker 1
```

To display statistics per VDOM:

```
# get test proxyworker 4444
```

To restart the proxy:

```
# get test proxyworker 99
```

How to examine the firewall session list

The firewall session list displays all sessions that the FortiGate has open. You can see if there are strange patterns, such as no sessions apart from the internal network, or all sessions are only to one IP address.

When you examine the firewall session list in the CLI, you can use filters to reduce the output. In the GUI, the filters are part of the interface.

To examine the firewall session list - GUI

Go to **FortiView > All Sessions**.

To examine the firewall session list - CLI

When you examine the firewall session list, there may be too many sessions to display. In this case, it's necessary to limit or filter the sessions displayed by source or destination address, or NAT'd address or port. If you want to filter by more than one of these, you need to enter a separate line for each value.

The following example shows filtering the session list based on a source address of 10.11.101.112:


```
FGT# diagnose system session filter src 10.11.101.112
FGT# diagnose system session list
```

The following example shows filtering the session list based on a destination address of 172.20.120.222.

```
FGT# diagnose system session filter dst 172.20.120.222
FGT# diagnose system session list
```

To clear all sessions corresponding to a filter - CLI

```
FGT# diagnose system session filter dst 172.20.120.222
FGT# diagnose system session clear
```

Check source NAT information

When you troubleshoot connections, remember NAT. NAT is especially important if you're troubleshooting from the remote end of the connection outside the firewall. On the **FortiView > All Sessions** list, pay attention to **NAT Source**, and **NAT Source Port**. These columns display the IP and port values after NAT has been applied.

Checking the NAT values can help you to ensure that they are the values you expect, and to ensure the remote end of the sessions can see the expected IP address and port number.

When you display the session list in the CLI, you can match the NAT'd source address (`nsrc`) and port (`nport`). This can be useful if multiple internal IP addresses are NAT'd to a common external-facing source IP address.

```
FGT# diagnose system session filter nsrc 172.20.120.122
FGT# diagnose system session filter nport 8888
FGT# diagnose system session list
```

How to check wireless information

Wireless connections, stations, and interfaces have different issues than other physical interfaces.

Troubleshooting station connection issue

To check whether a station entry is created on access control, use the following command:

```
FG600B3909600253 # diagnose wireless-controller wlaac -d sta
* vf=0 wtp=70 rId=2 wlan=open ip=0.0.0.0 mac=00:09:0f:db:c4:03 rssi=0 idle=148 bw=0 use=2
vf=0 wtp=70 rId=2 wlan=open ip=172.30.32.122 mac=00:25:9c:e0:47:88 rssi=-40 idle=0 bw=9
use=2
```

Enable diagnostics for a particular station

This example uses the station MAC address to find where it's failing:

```
FG600B3909600253 # diagnose wireless-controller wlaac sta_filter 00:25:9c:e0:47:88 1
Set filter sta 00:25:9c:e0:47:88 level 1
FG600B3909600253 # 71419.245 <ih> IEEE 802.11 mgmt::disassoc <== 00:25:9c:e0:47:88 vap
open rId 1 wId 0 00:09:0f:db:c4:03
71419.246 <dc> STA del 00:25:9c:e0:47:88 vap open rId 1 wId 0
71419.246 <cc> STA_CFG_REQ(34) sta 00:25:9c:e0:47:88 del ==> ws (0-192.168.35.1:5246) rId
1 wId 0
71419.246 <cc> STA del 00:25:9c:e0:47:88 vap open ws (0-192.168.35.1:5246) rId 1 wId 0
00:09:0f:db:c4:03 sec open reason I2C_STA_DEL
71419.247 <cc> STA_CFG_RESP(34) 00:25:9c:e0:47:88 <== ws (0-192.168.35.1:5246) rc 0
(Success).
```

How to verify connectivity to FortiGuard

You can verify connectivity to FortiGuard in the **Licenses** widget on the FortiGate **Dashboard**. When your FortiGate is connected to FortiGuard, a green check mark appears for available FortiGuard services.

To verify connectivity from the CLI, use the `execute ping service.fortiguard.net` and `execute ping update.fortiguard.net` commands.

Sample output:

```
FG100D# execute ping service.fortiguard.net
PING guard.fortinet.net (208.91.112.196): 56 data bytes
64 bytes from 208.91.112.196: icmp_seq=0 ttl=51 time=61.0 ms
64 bytes from 208.91.112.196: icmp_seq=1 ttl=51 time=60.0 ms
64 bytes from 208.91.112.196: icmp_seq=2 ttl=51 time=59.6 ms
64 bytes from 208.91.112.196: icmp_seq=3 ttl=51 time=58.9 ms
64 bytes from 208.91.112.196: icmp_seq=4 ttl=51 time=59.2 ms

--- guard.fortinet.net ping statistics ---
5 packets transmitted, 5 packets received, 0% packet loss
round-trip min/avg/max = 58.9/59.7/61.0 ms

FG100D# execute ping update.fortiguard.net
PING fds1.fortinet.com (208.91.112.68): 56 data bytes
64 bytes from 208.91.112.68: icmp_seq=0 ttl=53 time=62.0 ms
64 bytes from 208.91.112.68: icmp_seq=1 ttl=53 time=61.8 ms
64 bytes from 208.91.112.68: icmp_seq=2 ttl=53 time=61.3 ms
64 bytes from 208.91.112.68: icmp_seq=3 ttl=53 time=61.9 ms
64 bytes from 208.91.112.68: icmp_seq=4 ttl=53 time=61.8 ms
```

How to perform a sniffer trace (CLI and packet capture)

When you troubleshoot networks and routing in particular, it helps to look inside the headers of packets to determine if they're traveling along the route that you expect them to take. Packet sniffing can also be called a network tap, packet capture, or logic analyzing.



If your FortiGate has NP2, NP4, or NP6 interfaces that are offloading traffic, this will change the sniffer trace. Before you perform a trace on any NP2, NP4, or NP6 interfaces, you should disable offloading on those interfaces.

How do you sniff packets

Before you start sniffing packets on the CLI, you should prepare to capture the output to a file. A large amount of data may scroll by and you won't be able to see it without first saving it to a file. One method is to use a terminal program like `puTTY` to connect to the FortiGate CLI. Once the packet sniffing count is reached, you can end the session and analyze the output in the file.

The general form of the internal FortiOS packet sniffer command is:

```
diagnose sniffer packet <interface_name> <'filter'> <verbose> <count> <tsformat>
```

To stop the sniffer, type `CTRL+C`.

<interface_name>	The name of the interface to sniff, such as “port1” or “internal”. This can also be “any” to sniff all interfaces.
<'filter'>	What to look for in the information the sniffer reads. “none” indicates no filtering, and all packets are displayed as the other arguments indicate. The filter must be inside single quotes (').
<verbose>	The level of verbosity as one of: 1 - print header of packets 2 - print header and data from IP of packets 3 - print header and data from Ethernet of packets 4 - print header of packets with interface name
<count>	The number of packets the sniffer reads before stopping. If you don't put a number here, the sniffer will run until you stop it with <CTRL+C>.
<tsformat>	The format of timestamp. <ul style="list-style-type: none"> • a: absolute UTC time, yyyy-mm-dd hh:mm:ss.ms • l: absolute LOCAL time, yyyy-mm-dd hh:mm:ss.ms • otherwise: relative to the start of sniffing, ss.ms

For a simple sniffing example, enter the CLI command `diagnose sniffer packet port1 none 1 3`. This displays the next three packets on the port1 interface using no filtering, and verbose level 1. At this verbosity level, you can see the source IP and port, the destination IP and port, action (such as ack), and sequence numbers.

In the output below, port 443 indicates these are HTTPS packets and that 172.20.120.17 is both sending and receiving traffic.

```
Head_Office_620b # diagnose sniffer packet port1 none 1 3
interfaces=[port1]
filters=[none]

0.545306 172.20.120.17.52989 -> 172.20.120.141.443: psh 3177924955 ack 1854307757

0.545963 172.20.120.141.443 -> 172.20.120.17.52989: psh 1854307757 ack 3177925808

0.562409 172.20.120.17.52988 -> 172.20.120.141.443: psh 4225311614 ack 3314279933
```

For a more advanced example of packet sniffing, the following commands will report packets on any interface that are traveling between a computer with the host name of “PC1” and a computer with the host name of “PC2”. With verbosity 4 and above, the sniffer trace displays the interface names where traffic enters or leaves the FortiGate unit. Remember to stop the sniffer, type CTRL+C.

```
FGT# diagnose sniffer packet any "host <PC1> or host <PC2>" 4

or

FGT# diagnose sniffer packet any "(host <PC1> or host <PC2>) and icmp" 4
```

The following CLI command for a sniffer includes the ARP protocol in the filter which may be useful to troubleshoot a failure in the ARP resolution (for example, PC2 may be down and not responding to the FortiGate ARP requests).

```
FGT# diagnose sniffer packet any "host <PC1> or host <PC2> or arp" 4
```

How do you use packet capture

To use packet capture, the FortiGate must have a disk. You can enable the `capture-packet` in the firewall policy, using the following CLI commands:

```
config firewall policy
  edit <id>
    set capture-packet enable
  end
```

To configure packet capture filters, go to **Network > Packet Capture**.

When you add a packet capture filter, enter the following information and select **OK**.

Interface	<p>Select the interface to sniff from the drop-down menu.</p> <p>You must select one interface. You can't change the interface without deleting the filter and creating a new one, unlike the other fields.</p>
Max Packets to Save	<p>Enter the number of packets to capture before the filter stops.</p> <p>This number can't be zero. You can halt the capturing before this number is reached.</p>
Enable Filters	Select this option to specify filter fields
Host(s)	<p>Enter the IP address of one or more hosts</p> <p>Separate multiple hosts with commas. To enter a range, use a dash without spaces, for example 172.16.1.5-172.16.1.15, or enter a subnet.</p>
Port(s)	<p>Enter one or more ports to capture on the selected interface.</p> <p>Separate multiple ports with commas. To enter a range, use a dash without spaces, for example 88-90</p>
VLAN(s)	<p>Enter one or more VLANs (if any).</p> <p>Separate multiple VLANs with commas.</p>
Protocol	Enter one or more protocols. Separate multiple protocols with commas. To enter a range, use a dash without spaces, for example 1-6, 17, 21-25.
Include IPv6 Packets	Select this option if you're troubleshooting IPv6 networking, or if your network uses IPv6. Otherwise, leave it disabled.

Include Non-IP Packets

The protocols in the list are all IP based except for ICMP (ping). To capture non-IP based packets, select this feature. Examples of non-IP packets include IPsec, IGMP, ARP, and ICMP.

If you select a filter, you have the option to start and stop packet capture in the edit window, or download the captured packets. You can also see the filter status and the number of packets captured.

You can select the filter and start capturing packets. When the filter is running, the number of captured packets increases until it reaches the **Max Packet Count** or you stop it. When the filter is running, you can't download the output file.

When the packet capture is complete, you can download the *.pcap file. You must use a third party application, such as Wireshark, to read *.pcap files. This tool provides you with extensive analytics and the full contents of the packets that were captured.

To start, stop, or resume packet capture, use the symbols on the screen. These symbols are the same as those used for audio or video playback. Hover over the symbol to reveal explanatory text. Similarly, to download the *.pcap file, use the download symbol on the screen.

For more information on troubleshooting with packet capture and packet sniffing, see ["Packet sniffing and packet capture" on page 2894](#).

How to debug the packet flow

Traffic should come in and leave the FortiGate. If you determine that network traffic isn't entering and leaving the FortiGate as expected, debug the packet flow.

You can only perform debugging using CLI commands. Debugging the packet flow requires that you enter a number of debug commands as each one configures part of the debug action, with the final command starting the debug.



If your FortiGate has FortiASIC NP4 or NP6 interface pairs that are offloading traffic, this changes the packet flow. Before you perform the debug on any NP4 or NP6 interfaces, you should disable offloading on those interfaces. Enter `diagnose npu <interface pair> fastpath disable`, where `interface pair` is `np4`, `np6`, `np4lite`, or `np6lite`.

The following configuration assumes that PC1 is connected to the internal interface of the FortiGate and has an IP address of 10.11.101.200. PC1 is the host name of the computer.

To debug the packet flow in the CLI, enter the following commands:

```
FGT# diagnose debug disable
FGT# diagnose debug flow filter add <PC1>
FGT# diagnose debug flow show console enable
FGT# diagnose debug flow show function-name enable
FGT# diagnose debug flow trace start 100
FGT# diagnose debug enable
```

The `start 100` argument in the above list of commands limits the output to 100 packets from the flow. This is useful to look at the flow without flooding your log or displaying too much information.

To stop all other debug activities, enter the command:

```
FGT# diagnose debug flow trace stop
```

The following is an example of debug flow output for traffic that has no matching security policy, and is in turn blocked by the FortiGate unit. The denied message indicates that the traffic was blocked.

```
id=20085 trace_id=319 func=resolve_ip_tuple_fast line=2825 msg="vd-root received a packet
(proto=6, 192.168.129.136:2854->192.168.96.153:1863) from port3."

id=20085 trace_id=319 func=resolve_ip_tuple line=2924 msg="allocate a new session-
013004ac"

id=20085 trace_id=319 func=vf_ip4_route_input line=1597 msg="find a route: gw-
192.168.150.129 via port1"

id=20085 trace_id=319 func=fw_forward_handler line=248 msg=" Denied by forward policy
check"
```

Troubleshooting resources

Fortinet provides customers with resources of valuable information about FortiOS technical issues, including:

Technical documentation

Installation, Administration, and Quick Start Guides, as well as other technical documents, are available online at the [Fortinet Document Library](https://docs.fortinet.com): <https://docs.fortinet.com>

Fortinet video library

The [Fortinet Video Library](https://video.fortinet.com) hosts a collection of video which provide valuable information about Fortinet products.

<https://video.fortinet.com>

Release notes

Issues that arise after the technical documentation has been published will often be listed in the [Release Notes](#).

To find these, go to the [Fortinet Document Library](#).

<https://docs.fortinet.com/fortigate/release-information>

Knowledge base

The [Fortinet Knowledge Base](#) provides access to a variety of articles, white papers, and other documentation that provides technical insight into a range of Fortinet products. The Knowledge Base is available online at:

<http://kb.fortinet.com>

Fortinet technical discussion forums

An [online technical forum](#) allows administrators to contribute to discussions about issues that relate to their Fortinet products. Searching the forum can help an administrator identify if an issue has been experienced by another user. You can access the support forums at: <https://forum.fortinet.com/>

Fortinet training services online campus

The [Fortinet Training Services Online Campus](https://www.fortinet.com/training.html) hosts a collection of tutorials and training materials which you can use to increase your knowledge of Fortinet products. <https://www.fortinet.com/training.html>

Fortinet customer support

You defined your problem, researched a solution, put together a plan to find the solution, and executed that plan. At this point, if the problem hasn't been solved, it's time to contact [Fortinet Support](#) for assistance. Prepare yourself by reading [How to work with Fortinet Support](#) on Fortinet's Cookbook site.

<https://support.fortinet.com>

<http://cookbook.fortinet.com/how-to-work-with-fortinet-support/>

Chapter 28 - Virtual Domains

- [VDMs in NAT mode](#): detailed explanations and examples for configuring VDOM features for a FortiGate in NAT/route mode.
- [VDMs in transparent mode](#): detailed explanations and examples for configuring VDOM features for a FortiGate in transparent mode.
- [Inter-VDOM routing](#): concepts and scenarios for inter-VDOM routing.
- [Troubleshooting VDMs](#): diagnostic and troubleshooting information for some potential VDOM issues.



Before you begin using this guide, take a moment to note the following:

- By default, most FortiGate devices support 10 VDOMs. Many FortiGate models support purchasing a license key to increase the maximum number
 - This guide uses a FortiGate with interfaces named port1 through port4 for examples and procedures. The interface names on some models will vary. Where possible aliases for these ports are indicated to show their intended purpose and to help you determine which ports to use if your ports are labeled differently.
 - Administrators are assumed to be super_admin administrators unless otherwise specified. Some restrictions will apply to other administrators.
-

What's new in FortiOS 6.0

The following list contains new VDOM features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- ["Global security profiles" on page 2981](#)



Security profiles can still be created at the VDOM level, with that profile only available for that specific VDOM.

VDOMs overview

You can use virtual domains (VDOMs) to divide a FortiGate into two or more virtual devices that function independently. For each separate VDOM, you can create different configurations, including firewall policies, routing, VPNs, and security profiles.

Once you have created a VDOM, the steps you need to take to configure it are typically the same as if you were configuring a single FortiGate with VDOMs disabled. This chapter focuses on considerations that are unique to a network using VDOMs.

This section includes:

- [Benefits of virtual domains](#)
- [Enabling and accessing VDOMs](#)
- [Configuring additional VDOMs](#)



The Fortinet Security Fabric doesn't support FortiGate units with multiple VDOMs.

Benefits of virtual domains

VDOMs provide the following benefits:

- [Savings in physical space and power](#)
- [Using both NAT mode and transparent mode](#)
- [MSSP configuration](#)
- [Virtual clustering](#)

Savings in physical space and power

To increase the number of physical FortiGate devices, you require rack space, cables, and power. You also need to change your network configuration to accommodate the new devices. Finally, if you don't need as many devices in the future, you're left with expensive hardware that you aren't using.

Increasing the number of VDOMs requires no additional hardware and minimal changes to existing networking configurations. VDOMs save physical space and power. You're limited only by the size of your VDOM license and the physical resources of your FortiGate. By default, most FortiGate devices support a maximum of ten VDOMs, and many models allow you to buy a license to increase the maximum number of VDOMs.

Each physical FortiGate also requires a separate FortiGuard license to access security updates. VDOMs don't require you to buy separate licenses, as the same license is shared for all VDOMs located on the same FortiGate. When you update or upgrade the license, the changes are immediately available for all VDOMs.

Using both NAT mode and transparent mode

With multiple VDOMs, you can configure one VDOM configured in transparent mode and other VDOMs in NAT mode. In this configuration, you can use the transparent mode VDOM for troubleshooting your network and the NAT mode VDOMs for networking.

MSSP configuration

If you require a managed security service provider (MSSP) configuration, you can use VDOMs to provide a multi tenant solution, with each tenant's network connected to a unique VDOM that's configured to meet the network requirements. For each VDOM, you can either manage it globally using the management VDOM or allow tenants to manage their own VDOM.

Virtual clustering

Virtual clustering is an extension of FortiGate high availability for a cluster of two FortiGate units with multiple VDOMs. Virtual clustering provides failover protection for a multiple VDOM configuration and can load balance traffic between the VDOMs to improve overall network performance.

Virtual clustering load balancing efficiently load balances all traffic between VDOMs (including TCP and UDP traffic, UTM traffic and VoIP traffic) and can be adjusted in real time to actively optimize load sharing between the cluster units without affecting the smooth operation of the VDOMs in the cluster.

Enabling and accessing VDOMs

While a FortiGate VDOM is essentially the same as a physical FortiGate, there are some small differences. After you enable VDOMs on your FortiGate, you should take time to familiarize yourself with the changes.

This section includes:

- [The root VDOM](#)
- [Enabling virtual domains](#)
- [Global and per-VDOM settings](#)
- [Changes to the GUI and CLI](#)
- [Enabling and accessing VDOMs](#)
- [Resource settings](#)
- [Increasing the maximum number of VDOMs](#)

The root VDOM

On every FortiGate there is a root VDOM that you can't delete, whether or not VDOMs are enabled. When VDOMs are disabled, the root VDOM isn't visible because it doesn't need to be. When VDOMs are enabled, the root VDOM is visible. The root VDOM is the only VDOM available for configuration, until you

Typically, you use the root VDOM as the management VDOM. By connecting to the management VDOM, you can access the global settings for the FortiGate as well as the settings for each individual VDOM. You can set any VDOM to be the management VDOM.

Enabling virtual domains

VDOMs are disabled by default. To enable VDOMs, use the following instructions.

To enable VDOMs - GUI:

1. Go to **System > Settings**.
2. Under **Operations Settings**, enable **Virtual Domains**.

The FortiGate logs off all sessions. You can now log in again as admin.

To enable VDOMs - CLI:

```
config system global
    set vdom-admin enable
end
```

Global and per-VDOM settings

Settings that you configure outside a VDOM are called global settings. These settings affect the entire FortiGate and include areas such as interfaces, HA, maintenance, some antivirus, and some logging. In general, global settings are settings that should only be changed by your top level administrator.

Settings that you configure within a VDOM are called VDOM settings. These settings affect only a specific VDOM and include areas such as operating mode, routing, firewall, VPN, some antivirus, some logging settings, and reporting.

When virtual domains aren't enabled, the entire FortiGate is effectively a single VDOM, but per-VDOM limits apply. For some resource types, the global limit can't be reached with only one VDOM.

Changes to the GUI and CLI

When you enable VDOMs, the FortiGate GUI and the CLI change, allowing you to manage both global settings and per-VDOM settings. Only admin accounts using the super_admin profiles can access global options and settings for all VDOMs. Other administrator accounts can configure only the VDOM they're assigned to.

Other changes only affect either the GUI or the CLI.

GUI:

- When you access the management VDOM (this is the root VDOM by default), you can use the drop-down menu in the top left of the GUI to switch between global and per-VDOM settings. Some menu items only appear under **Global**, while others only appear as per-VDOM settings.
- A menu item is available at **Global > System > VDOM**. You use this to create, edit, and delete VDOMs.
- A menu item is available at **Global > System > Global Resources**. You use this to manage how system resources are shared between VDOMs.

CLI:

- To configure global settings, you must first enter the following CLI to select global options:

```
config global
```

- To configure per-VDOM settings, you must first enter the following CLI to select a VDOM:

```
config vdom
    edit <vdom_name>
```

Resource settings

Your FortiGate has a limited amount of hardware resources, such as memory, disk storage, and CPU operations. When you use VDOMs, you can control how resources are shared between each VDOM to optimize resource usage. This allows you to ensure the proper level of service is maintained on each VDOM.

For example, if one VDOM is connected to a web server and logging server, and a second VDOM is connected to an internal network of 20 users, these two VDOMs require different levels of resources. The first VDOM requires many sessions but no user accounts. The second VDOM is the opposite and requires user accounts and management resources, but fewer sessions.

Global resources

Global resources apply to the entire FortiGate. By default, the values are set to their maximum values. These values vary by FortiGate model because each model has different hardware capabilities. It can be useful to change the maximum values for some resources to ensure there is enough memory available for other resources that may be more important to your configuration.

To use the earlier example, if your FortiGate is protecting a number of web servers and other publicly accessible servers, you should maximize the available sessions and proxies and minimize other settings that are unused, such as user settings, VPNs, and dial-up tunnels.

To view the resource list, go to **Global > System > Global Resources**. You can also use the following CLI command:

```
config global
  config system resource-limits
  get
```

Note that some global resources are only visible if your FortiGate supports those resources. For example, the quota for logging to disk is only visible when your FortiGate has a hard disk.

For explicit proxies, when you configure limits on the number of concurrent users, you need to allow for the number of users based on their authentication method. Otherwise you may run out of user resources.



Each session-based authenticated user is counted as a single user using their authentication membership (RADIUS, LDAP, FSAE, local database, etc.) to match users in other sessions. So one authenticated user in multiple sessions is still one user.

For all other situations, the source IP address is used to determine a user. All sessions from a single source address are assumed to be from the same user.

Per-VDOM resource settings

While global resources apply to resources shared by the entire FortiGate, per-VDOM resources are specific to only one VDOM. By default, all per-VDOM resource settings are set to allow the maximum.

Each VDOM has its own resource settings, including both maximum and minimum levels. The maximum level is the highest amount of that resource that the VDOM can use if it is available on the FortiGate. Minimum levels are guaranteed levels that are always available, no matter what resources the other VDOMs are using.

For example, one VDOM, called VDOM-1, has a maximum of 5000 sessions and a minimum of 1000 sessions. If the FortiGate has a global maximum of 20,000 sessions split among 10 VDOMs, it is possible that VDOM-1 won't be able to reach the 5000 session maximum. However, at all times VDOM-1 is guaranteed to have 1000 sessions available.

To view per-VDOM resource settings - GUI:

1. Select **Global > System > VDOM**.
2. Select the root VDOM, and select **Edit**.
3. Adjust the settings in the **Resource Usage** section of the page.
4. Select **OK**.

To view per-VDOM resource settings - CLI:

```
config global
  config system vdom-property
    edit root
    get
```

Increasing the maximum number of VDOMs

By default, most FortiGate models support configuring a maximum of 10 VDOMs. For certain models, you can purchase a license key to increase the maximum number of VDOMs.

To find out how many VDOMs your FortiGate can support, refer to the data sheet for your model. For more information, see the [Fortinet Data Sheets](#).



It is important to back up your configuration before upgrading the VDOM license on your FortiGate, especially if you're using HA mode.

To obtain a VDOM license key

1. Log in with a super_admin account.
2. Go to the **Dashboard**.
3. Record your FortiGate serial number as shown in **System Information** widget.
4. In the **License Information** widget, locate **Virtual Domain** and select **Purchase More**.



If you don't see the **Purchase More** option, your FortiGate model does not support more than 10 VDOMs.

5. You are taken to the [Fortinet Support website](#) where you can log in and purchase a license key.
6. After you receive your license key, go to the Dashboard and select **Upload License** under **License Information, Virtual Domains**.
7. In the **Input License Key** field, enter the 32-character license key you received from Fortinet Support.
8. Select **Apply**.

To verify the new VDOM license, in global configuration go to **System > Dashboard**. The **Licenses** widget shows the current number and total allowed number of VDOMs.

Configuring additional VDOMs

After enabling VDOMs, you can create additional VDOMs on your FortiGate and configure them for your network requirements.

This section includes:

- [Creating a VDOM](#)
- [Changing the management VDOM](#)
- [Adding interfaces to a VDOM](#)
- [Per-VDOM administrators](#)
- [Certificate management](#)
- [Disabling a VDOM](#)
- [Deleting a VDOM](#)

Creating a VDOM



You may experience reduced FortiGate performance if you create a large number of VDOMs.

To create new VDOMs, you must have a super_admin profile account. You must use the management VDOM (the root VDOM, by default) to create new VDOMs.

By default, new VDOMs are set to NAT/route operation mode. If you want a VDOM to be in transparent operation mode, you must manually change it using the CLI. For more information, see the [FortiOS CLI Reference](#).

To create a VDOM - GUI:

1. Connect to the management VDOM.
2. Go to **Global > System > VDOM** and select **Create New**.
3. Enter a unique **Name**. VDOM names have the following restrictions:
 - Only letters, numbers, “-”, and “_” are allowed
 - No more than 11 characters are allowed
 - No spaces are allowed
 - VDOMs can't have the same names as interfaces, zones, switch interfaces, or other VDOMs
4. Enter a short and descriptive comment to identify this VDOM.
5. Select **OK**.

To create a VDOM - CLI:

```
config system vdom
  edit <new_vdom_name>
end
```



If you want to edit an existing VDOM in the CLI and mistype the name, a new VDOM is created with this name.

The new VDOM can either be renamed or deleted. For more information, see [Deleting a VDOM](#).

Changing the management VDOM



You can't change the management VDOM if any administrators are using RADIUS authentication.

Once you have two or more VDOMs, you can change the management VDOM. The management VDOM must have Internet access.

You use the management VDOM to access global settings on the FortiGate, as well as for the following services:

- DNS lookups
- Logging to a FortiAnalyzer or syslog
- FortiGuard service
- Sending alert emails
- Network time protocol traffic (NTP)
- Sending SNMP traps
- Quarantining suspicious files and email

To change the management VDOM - GUI:

1. Select **Global > System > VDOM**.
2. Select the new management VDOM.
3. Select **Switch Management**.
4. Select **OK** to confirm the change.

To change the management VDOM - CLI:

```
config global
  config system global
    set management-vdom <vdom_name>
  end
end
```

Adding interfaces to a VDOM

Once you create a VDOM, you can add network interfaces to it. You can only assign an interface to a single VDOM. By default, all interfaces are assigned to the root VDOM.

If the existing configuration references an interface, you won't be able to change the VDOM assignment for that interface. Because some FortiGate models have a default configuration, you may need to delete existing policies and routes in order to add a particular interface to a new VDOM.

To add an interface to a VDOM - GUI:

1. Connect to the management VDOM.
2. Go to **Global > Network > Interfaces** and edit the interface.
3. Set **Virtual Domain** to the appropriate VDOM.
4. Select **OK**.

To add an interface to a VDOM - CLI:

```
config global
    config system interface
        edit <interface_name>
            set vdom <VDOM_name>
        next
    end
end
```

Per-VDOM administrators

After you enable VDOMs, you can create administrators with access to several VDOMs or limited to a single VDOM, called per-VDOM administrators.

Per-VDOM administrators must have either the `prof_admin` profile or a custom profile. Administrators who have the `super_admin` profile have access to all VDOMs on the FortiGate. For more information about administrator profiles, see the *System Administration handbook*.

Per-VDOM administrators must access the FortiGate through network interfaces that belong to those VDOMs, which must be configured to allow management access. The administrator can also connect using the console interface.

When per-VDOM administrators log into their virtual domain, they see a different dashboard than the global administrator sees. The VDOM dashboard displays information only relevant to that VDOM, while information about global settings or other VDOMs aren't shown.

Information	Per-VDOM	Global
System information	read-only	yes
License information	no	yes
CLI console	yes	yes
Unit operation	read-only	yes
Alert message console	no	yes
Top sessions	limited to VDOM sessions	yes
Traffic	limited to VDOM interfaces	yes
Statistics	yes	yes

You can create administrators globally or per-VDOM. In order to assign an administrator to multiple VDOMs, you must create the account at the global level. The steps for creating an administrator at the global level are shown below.

To create an administrator at the per-VDOM level, you use the same procedure used to create administrators on a FortiGate where VDOMs are disabled. However, the `super_admin` profile can't be used.

To create per-VDOM administrators - GUI:

1. Connect to the management VDOM.
2. Go to **Global > System > Administrators** and select **Create New**.
3. Set the **User Name** for the account.
4. Set and confirm the **Password**.
5. Set **Type** to **Local User**.
6. Remove the **root** VDOM from the **Virtual Domains** list, then add the appropriate VDOM.
7. Select **OK**.

To create per-VDOM administrators - CLI:

```
config global
  config system admin
    edit <name>
      set vdom <VDOM_name>
      set password <password>
      set accprofile <admin_profile>
      ...
    next
  end
end
```

Certificate management

Certificates can be uploaded to the FortiGate certificate store as either global or per-VDOM. Any certificates uploaded to the global certificate store will be available to all VDOMs on the FortiGate. A certificate uploaded to a VDOM-specific certificate store are only available for that VDOM.

The following factory default certificates are unique to each VDOM and are automatically generated when a new VDOM is added:

- Fortinet_CA_SSL
- Fortinet_SSL
- PositiveSSL_CA
- Fortinet_Wifi
- Fortinet_Factory

Disabling a VDOM

When you create a new VDOM, it's enabled by default. You can configure a VDOM only when it is enabled. You must enable the management VDOM.

Disabled VDOMs are considered offline. The configuration remains, but you can't use the VDOM and only the `super_admin` administrator can view it. You can assign interfaces to a disabled VDOM.

To disable a VDOM - GUI:

1. Go to **Global > System > VDOM**.
2. Open the VDOM for editing.
3. Ensure **Enable** is not selected.
4. Select **OK**.

To disable a VDOM - CLI:

```
config vdom
  edit <name>
    config system settings
      set status disable
    end
  end
```

Deleting a VDOM

You can only delete VDOMs that are enabled.

Deleting a VDOM removes it from the FortiGate configuration. You can't delete the root VDOM or the management VDOM.

You can only delete VDOMs that aren't referenced by the current configuration, including any per-VDOM objects. Before you delete a VDOM, check for, and remove the following objects that refer to that VDOM or its components:

- Routing - both static and dynamic routes
- Firewall addresses, policies, groups, or other settings
- Security profiles
- VPN configuration
- Users or user groups
- Logging
- DHCP servers
- Network interfaces, zones, and custom DNS servers
- VDOM administrators

Before you delete a VDOM, it's recommended that you re-assign interfaces assigned to that VDOM to the root VDOM.

To delete a VDOM - GUI:

1. Go to **Global > System > VDOM**.
2. Select the check box for the VDOM and then select the **Delete** icon.
3. Confirm the deletion.

To delete a VDOM - CLI:

```
config vdom
```

```
delete test-vdom
end
```

VDOMs in NAT mode

By default, VDOMs operate in NAT mode. In this mode, you install the VDOM as a gateway or router between two networks, typically a private network and the Internet. In this configuration, the VDOM uses network address translation (NAT) to hide the private IP addresses of network devices.

You can use VDOMs in NAT mode and transparent mode together on the same FortiGate. For more information about transparent mode, see ["VDOMs in transparent mode" on page 2964](#).

This chapter contains the following sections:

- [Using a VDOM in NAT/route mode](#)
- [Example configuration: VDOM in NAT/route mode](#)

Using a VDOM in NAT/route mode

This section contains information about how to configure a VDOM in NAT/route mode, including the following:

- [Configuring VDOM routing](#)
- [Configuring security policies](#)
- [Changing the inspection mode](#)
- [Using a VDOM in NAT/route mode](#)
- [Configuring VPNs for a VDOM](#)
- [Using a VDOM in NAT/route mode](#)

Configuring VDOM routing

Routing is VDOM-specific. Each VDOM should have a default static route configured as a minimum. Within a VDOM, routing is the same as routing on your FortiGate without VDOMs enabled.

When configuring dynamic routing on a VDOM, other VDOMs on the FortiGate can be neighbors. The following topics give a brief introduction to the routing protocols, and show specific examples of how to configure dynamic routing for VDOMs. Figures are included to show the FortiGate configuration after the successful completion of the routing example.

Default static route for a VDOM

The routing you define applies only to network traffic entering non-ssl interfaces belonging to this VDOM. Set the administrative distance high enough, typically 20, so that automatically configured routes will be preferred to the default.

In the following procedure, it is assumed that a VDOM called "Client2" exists. The procedure will create a default static route for this VDOM. The route has a destination IP of 0.0.0.0, on the port3 interface. It has a gateway of 10.10.10.1, and an administrative distance of 20.

The values used in this procedure are very standard, and this procedure should be part of configuring all VDOMs.

To add a default static route for a VDOM - GUI:

1. In **Virtual Domains**, select the client2 VDOM.
2. Go to **Network > Static Routes**.
3. Select **Create New**.
4. Enter the following information and select **OK**:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port2
Gateway	10.10.10.1
Distance	20

To add a default static route for a VDOM - CLI:

```
config vdom
  edit client2
    config router static
      edit 4
        set device port2
        set dst 0.0.0.0 0.0.0.0
        set gateway 10.10.10.1
        set distance 20
      end
    end
  end
```

Dynamic routing in VDOMs

Dynamic routing is VDOM-specific, like all other routing. Dynamic routing configuration is the same with VDOMs as with your FortiGate without VDOMs enabled, once you're at the routing menu. If you have multiple VDOMs configured, the dynamic routing configuration between them can become quite complex.

VDOMs provide some interesting changes to dynamic routing. Each VDOM can be a neighbor to the other VDOMs. This is useful in simulating a dynamic routing area or AS or network using only your FortiGate.

You can separate different types of routing to different VDOMs if required. This allows for easier troubleshooting. This is very useful if your FortiGate is on the border of a number of different routing domains.

For more information on dynamic routing in FortiOS, see the *Networking Handbook*.

Inter-VDOM links must have IP addresses assigned to them if they are part of a dynamic routing configuration. Inter-VDOM links may or may not have IP addresses assigned to them. Without IP addresses, you need to be careful how you configure routing. While the default static route can be assigned an address of 0.0.0.0 and rely instead on the interface, dynamic routing almost always requires an IP address.

RIP

The RIP dynamic routing protocol uses hop count to determine the best route, with a hop count of 1 being directly attached to the interface and a hop count of 16 being unreachable. For example if two VDOMs on the same FortiGate are RIP neighbors, they have a hop count of 1.

OSPF

OSPF communicates the status of its network links to adjacent neighbor routers instead of the complete routing table. When compared to RIP, OSPF is more suitable for large networks, it is not limited by hop count, and is more complex to configure. For smaller OSPF configurations its easiest to just use the backbone area, instead of multiple areas.

BGP

BGP is an Internet gateway protocol (IGP) used to connect autonomous systems (ASes) and is used by Internet service providers (ISPs). BGP stores the full path, or path vector, to a destination and its attributes which aid in proper routing.

Configuring security policies

Security policies are VDOM-specific. This means that all firewall settings for a VDOM, such as firewall addresses and security policies, are configured within the VDOM.

In VDOMs, all firewall related objects are configured per-VDOM including addresses, service groups, security profiles, schedules, traffic shaping, and so on. If you want firewall addresses, you will have to create them on each VDOM separately. If you have many addresses, and VDOMs this can be tedious and time consuming. Consider using a FortiManager unit to manage your VDOM configuration — it can get firewall objects from a configured VDOM or FortiGate, and push those objects to many other VDOMs or FortiGate devices. See the [FortiManager Administration Guide](#).



You can customize the **Policy** display by including some or all columns, and customize the column order onscreen. Due to this feature, security policy screen shots may not appear the same as on your screen.

Configuring a security policy for a VDOM

Your security policies can involve only the interfaces, zones, and firewall addresses that are part of the current VDOM, and they are only visible when you're viewing the current VDOM. The security policies of this VDOM filter the network traffic on the interfaces and VLAN subinterfaces in this VDOM.

A firewall service group can be configured to group multiple services into one service group. When a descriptive name is used, service groups make it easier for an administrator to quickly determine what services are allowed by a security policy.

In the following procedure, it is assumed that a VDOM called `Client2` exists. The procedure will configure an outgoing security policy. The security policy will allow all HTTPS, SSH, and DNS traffic for the `SalesLocal` address group on `VLAN_200` going to all addresses on port3. This traffic will be scanned and logged.

To configure a security policy for a VDOM - GUI:

1. In **Virtual Domains**, select the `client2` VDOM.
2. Go to **Policy & Objects > IPv4 Policy**.
3. Select **Create New**.
4. Enter the following information and select **OK**:

Name	Client2-outgoing
Incoming Interface	VLAN_200
Outgoing Interface	port3
Source Address	SalesLocal
Destination Address	any
Schedule	always
Service	HTTPS, SSH, DNS
Action	ACCEPT
Log Allowed Traffic	enable

To configure a security policy for a VDOM - CLI:

```

config vdom
  edit Client2
    config firewall policy
      edit 12
        set srcintf VLAN_200
        set srcaddr SalesLocal
        set dstintf port3
        set dstaddr any
        set schedule always
        set service HTTPS SSH
        set action accept
        set status enable
        set logtraffic enable
      end
    end
  end

```

Changing the inspection mode

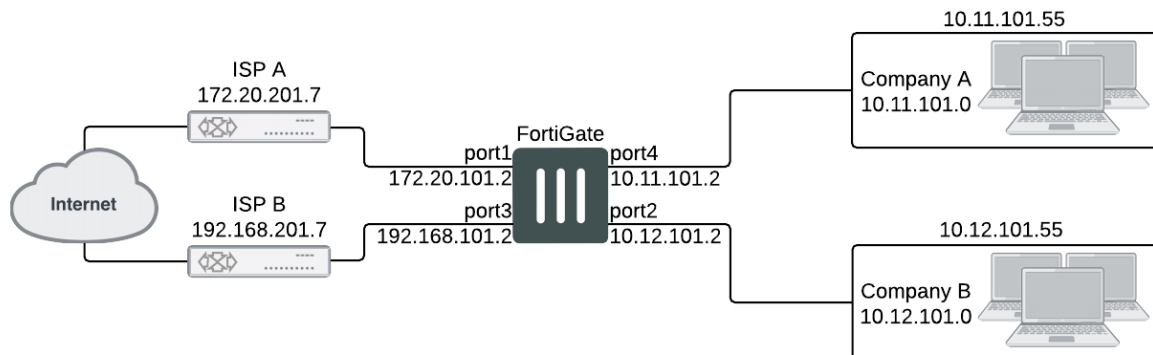
If you wish to change the inspection mode for a VDOM, go to **System > VDOM** and edit the VDOM you want to configure. Set **Inspection Mode** to either **Proxy** or **Flow-based**.

VDMs on the same FortiGate can use different inspection modes.

Configuring VPNs for a VDOM

Virtual Private Networking (VPN) settings are VDOM-specific, and must be configured within each VDOM. Configurations for IPsec Tunnel, IPsec Interface, PPTP and SSL are VDOM-specific.

Example configuration: VDOM in NAT/route mode



Company A and Company B each have their own internal networks and their own ISPs. They share a FortiGate that is configured with two separate VDOMs, with each VDOM running in NAT/route mode enabling separate configuration of network protection profiles. Each ISP is connected to a different interface on the FortiGate.

This network example was chosen to illustrate one of the most typical VDOM configurations.

This example has the following sections:

- [Network topology and assumptions](#)
- [General configuration steps](#)
- [Creating the VDOMs](#)
- [Configuring the FortiGate interfaces](#)
- [Configuring the vdomA VDOM](#)
- [Configuring the vdomB VDOM](#)
- [Testing the configuration](#)

Network topology and assumptions

Both companies have their own ISPs and their own internal interface, external interface, and VDOM on the FortiGate.

For easier configuration, the following IP addressing is used:

- all IP addresses on the FortiGate end in “.2” such as 10.11.101.2.
- all IP addresses for ISPs end in “.7”, such as 172.20.201.7.
- all internal networks are 10.*.* networks, and sample internal addresses end in “.55”.

The IP address matrix for this example is as follows.

Address	Company A	Company B
ISP	172.20.201.7	192.168.201.7
Internal network	10.11.101.0	10.012.101.0

Address	Company A	Company B
FortiGate / VDOM	172.20.201.2 (port1)	192.168.201.2 (port3)
	10.11.101.2 (port4)	10.012.101.2 (port2)

The Company A internal network is on the 10.11.101.0/255.255.255.0 subnet. The Company B internal network is on the 10.12.101.0/255.255.255.0 subnet.

There are no switches or routers required for this configuration.

There are no VLANs in this network topology.

The interfaces used in this example are port1 through port4. Different FortiGate models may have different interface labels. port1 and port3 are used as external interfaces. port2 and port4 are internal interfaces.

The administrator is a super_admin account. If you're using a non-super_admin account, refer to "Global and per-VDOM settings" to see which parts a non-super_admin account can also configure.

When configuring security policies in the CLI always choose a policy number that is higher than any existing policy numbers, select `services` before `profile-status`, and `profile-status` before `profile`. If these commands are not entered in that order, they may not be available to enter.

General configuration steps

For best results in this configuration, follow the procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. [Creating the VDOMs](#)
2. [Configuring the FortiGate interfaces](#)
3. [Configuring the vdomA VDOM](#), and [Configuring the vdomB VDOM](#)
4. [Testing the configuration](#)

Creating the VDOMs

In this example, two new VDOMs are created — vdomA for Company A and vdomB for Company B. These VDOMs will keep the traffic for these two companies separate while enabling each company to access its own ISP.

To create two VDOMs - GUI:

1. Log in with a super_admin account.
2. Go to **Global > System > VDOM**, and select **Create New**.
3. Enter `vdomA` and select **OK**.
4. Select **OK** again to return to the VDOM list.
5. Select **Create New**.
6. Enter `vdomB` and select **OK**.

To create two VDOMs - CLI:

```
config vdom
edit vdomA
next
```

```
edit vdomB
end
```

Configuring the FortiGate interfaces

This section configures the interfaces that connect to the companies' internal networks, and to the companies' ISPs.

All interfaces on the FortiGate will be configured with an IP address ending in ".2" such as 10.11.101.2. This will simplify network administration both for the companies, and for the FortiGate global administrator. Also the internal addresses for each company differ in the second octet of their IP address - Company A is 10.11.*, and Company B is 10.12.*.

This section includes the following topics:

- [Configuring the vdomA interfaces](#)
- [Configuring the vdomB interfaces](#)



If you can't change the VDOM of a network interface it is because something is referring to that interface that needs to be deleted. Once all the references are deleted the interface will be available to switch to a different VDOM. For example a common reference to the external interface is the default static route entry. See [Example configuration: VDOM in NAT/route mode](#).

Configuring the vdomA interfaces

The vdomA VDOM includes two FortiGate interfaces: port1 and external.

The port4 interface connects the Company A internal network to the FortiGate, and shares the internal network subnet of 10.11.101.0/255.255.255.0.

The external interface connects the FortiGate to ISP A and the Internet. It shares the ISP A subnet of 172.20.201.0/255.255.255.0.

To configure the vdomA interfaces - GUI:

1. Go to **Global > Network > Interfaces**.
2. Select **Edit** on the port1 interface.
3. Enter the following information and select **OK**:

Virtual Domain	vdomA
Addressing mode	Manual
IP/Netmask	172.20.201.2/255.255.255.0

4. Select **Edit** on the port4 interface.
5. Enter the following information and select **OK**:

Virtual Domain	vdomA
Addressing mode	Manual
IP/Netmask	10.11.101.2/255.255.255.0

To configure the vdomA interfaces - CLI:

```

config global
config system interface
edit port1
set vdom vdomA
set mode static
set ip 172.20.201.2 255.255.255.0
next
edit port4
set vdom ABCDomain
set mode static
set ip 10.11.101.2 255.255.255.0
end

```

Configuring the vdomB interfaces

The vdomB VDOM uses two FortiGate interfaces: port2 and port3.

The port2 interface connects the Company B internal network to the FortiGate, and shares the internal network subnet of 10.12.101.0/255.255.255.0.

The port3 interface connects the FortiGate to ISP B and the Internet. It shares the ISP B subnet of 192.168.201.0/255.255.255.0.

To configure the vdomB interfaces - GUI:

1. Go to **Global > Network > Interfaces**.
2. Select **Edit** on the port3 interface.
3. Enter the following information and select **OK**:

Virtual domain	vdomB
Addressing mode	Manual
IP/Netmask	192.168.201.2/255.255.255.0

4. Select **Edit** on the port2 interface.
5. Enter the following information and select **OK**:

Virtual domain	vdomB
Addressing mode	Manual
IP/Netmask	10.12.101.2/255.255.255.0

To configure the vdomB interfaces - CLI:

```

config global
config system interface
edit port3
set vdom vdomB
set mode static
set ip 192.168.201.2 255.255.255.0
next

```

```
edit port2
    set vdom vdomB
    set mode static
    set ip 10.12.101.2 255.255.255.0
end
```

Configuring the vdomA VDOM

With the VDOMs created and the ISPs connected, the next step is to configure the vdomA VDOM.

Configuring the vdomA includes the following:

- [Adding vdomA firewall addresses](#)
- [Adding the vdomA security policy](#)
- [Adding the vdomA default route](#)

Adding vdomA firewall addresses

You need to define the addresses used by Company A's internal network for use in security policies. This internal network is the 10.11.101.0/255.255.255.0 subnet.

The FortiGate provides one default address, "all", that you can use when a security policy applies to all addresses as the source or destination of a packet.

To add the vdomA firewall addresses - GUI:

1. In **Virtual Domains**, select **vdomA**.
2. Go to **Policy & Objects > Addresses**.
3. Select **Create New**.
4. Enter the following information and select **OK**:

Address Name	Ainternal
Type	Subnet / IP Range
Subnet / IP Range	10.11.101.0/255.255.255.0
Interface	port4

To add the ABCDomain VDOM firewall addresses - CLI:

```
config vdom
    edit vdomA
        config firewall address
            edit Ainternal
                set type ipmask
                set subnet 10.11.101.0 255.255.255.0
            end
        end
    end
```

Adding the vdomA security policy

You need to add the vdomA security policy to allow traffic from the internal network to reach the external network, and from the external network to internal as well. You need two policies for this domain.

To add the vdomA security policy - GUI:

1. In **Virtual Domains**, select **vdomA**.
2. Go to **Policy & Objects > IPv4 Policy**.
3. Select **Create New**.
4. Enter the following information and select **OK**:

Name	VDOMA-internal-to-external
Incoming Interface	port4
Outgoing Interface	port1
Source Address	Ainternal
Destination Address	all
Schedule	Always
Service	ANY
Action	ACCEPT

5. Select **Create New**.
6. Enter the following information and select **OK**:

Name	VDOMA-external-to-internal
Incoming Interface	port1
Outgoing Interface	port4
Source Address	all
Destination Address	Ainternal
Schedule	Always
Service	ANY
Action	ACCEPT

To add the vdomA security policy - CLI:

```

config vdom
  edit vdomA
    config firewall policy
      edit 1
        set srcintf port4
        set srcaddr Ainternal
        set dstintf port1
        set dstaddr all
        set schedule always
        set service ANY
        set action accept

```

```

        set status enable
    next
    edit 2
        set srcintf port1
        set srcaddr all
        set dstintf port4
        set dstaddr Ainternal
        set schedule always
        set service ANY
        set action accept
        set status enable
    end

```

Adding the vdomA default route

You also need to define a default route to direct packets from the Company A internal network to ISP A. Every VDOM needs a default static route, as a minimum, to handle traffic addressed to external networks such as the Internet.

The administrative distance should be set slightly higher than other routes. Lower admin distances will get checked first, and this default route will only be used as a last resort.

To add a default route to the vdomA - GUI:

1. For **Virtual Domains**, select **vdomA**.
2. Go to **Network > Static Routes**.
3. Select **Create New**.
4. Enter the following information and select **OK**:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port1
Gateway	172.20.201.7
Distance	20

To add a default route to the vdomA - CLI:

```

config vdom
    edit vdomA
        config router static
            edit 1
                set device port1
                set gateway 172.20.201.7
            end
        end
    end

```

Configuring the vdomB VDOM

In this example, the vdomB VDOM is used for Company B. Firewall and routing settings are specific to a single VDOM.

vdomB includes the FortiGate port2 interface to connect to the Company B internal network, and the FortiGate port3 interface to connect to ISP B. Security policies are needed to allow traffic from port2 to external and from external to port2 interfaces.

This section includes the following topics:

- [Adding the vdomB firewall address](#)
- [Adding the vdomB security policy](#)
- [Adding a default route to the vdomB VDOM](#)

Adding the vdomB firewall address

You need to define addresses for use in security policies. In this example, the vdomB VDOM needs an address for the port2 interface and the “all” address.

To add the vdomB firewall address - GUI:

1. In **Virtual Domains**, select **vdomB**.
2. Go to **Policy & Objects > Addresses**.
3. Select **Create New**.
4. Enter the following information and select **OK**:

Address Name	Binternal
Type	Subnet / IP Range
Subnet / IP Range	10.12.101.0/255.255.255.0
Interface	port2

To add the vdomB firewall address - CLI:

```
config vdom
  edit vdomB
    config firewall address
      edit Binternal
        set type ipmask
        set subnet 10.12.101.0 255.255.255.0
      end
    end
  end
```

Adding the vdomB security policy

You also need a security policy for the Company B domain. In this example, the security policy allows all traffic.

To add the vdomB security policy - GUI:

1. Log in with a super_admin account.
2. In **Virtual Domains**, select vdomB.
3. Go to **Policy & Objects > IPv4 Policy**
4. Select **Create New**.
5. Enter the following information and select **OK**:

Name	VDOMB-internal-to-external
Incoming Interface	port2
Outgoing Interface	port3
Source Address	Binternal
Destination Address	all
Schedule	Always
Service	ANY
Action	ACCEPT

6. Select **Create New**.
7. Enter the following information and select **OK**:

Name	VDOMB-external-to-internal
Incoming Interface	port3
Outgoing Interface	port2
Source Address	all
Destination Address	Binternal
Schedule	Always
Service	ANY
Action	ACCEPT

To add the vdomB security policy - CLI:

```

config vdom
  edit vdomB
    config firewall policy
      edit 1
        set srcintf port2
        set dstintf port3
        set srcaddr Binternal
        set dstaddr all
        set schedule always
        set service ANY
        set action accept
        set status enable
      edit 1
        set srcintf port3
        set dstintf port2
        set srcaddr all
        set dstaddr Binternal
        set schedule always
        set service ANY

```



```

        set action accept
        set status enable
    end
end

```

Adding a default route to the vdomB VDOM

You need to define a default route to direct packets to ISP B.

To add a default route to the vdomB VDOM - GUI:

1. Log in as the super_admin administrator.
2. In **Virtual Domains**, select vdomB.
3. Go to **Network > Static Routes**.
4. Select **Create New**.
5. Enter the following information and select **OK**:

Destination IP/Mask	0.0.0.0/0.0.0.0
Device	port3
Gateway	192.168.201.7
Distance	20

To add a default route to the vdomB VDOM - CLI:

```

config vdom
  edit vdomB
    config router static
      edit 1
        set dst 0.0.0.0/0
        set device external
        set gateway 192.168.201.7
      end
    end
  end
end

```

Testing the configuration

Once you have completed configuration for both company VDOMs, you can use diagnostic commands, such as `tracert` in Windows, to test traffic routed through the FortiGate. Alternately, you can use the `traceroute` command on a Linux system with similar output.

Possible errors during the traceroute test are:

- “* * * Request timed out” - the trace was not able to make the next connection towards the destination fast enough
- “Destination host unreachable” - after a number of timed-out responses the trace will give up

Possible reasons for these errors are bad connections or configuration errors.

For additional troubleshooting, see [Troubleshooting VDOMs](#).

Testing traffic from the internal network to the ISP

In this example, a route is traced from the Company A internal network to ISP A. The test was run on a Windows PC with an IP address of 10.11.101.55.

The output here indicates three hops between the source and destination, the IP address of each hop, and that the trace was successful.

From the Company A internal network, access a command prompt and enter this command:

```
C:\>tracert 172.20.201.7
Tracing route to 172.20.201.7 over a maximum of 30 hops:
  1  <10 ms  <10 ms  <10 ms  10.11.101.2
  2  <10 ms  <10 ms  <10 ms  172.20.201.2
  3  <10 ms  <10 ms  <10 ms  172.20.201.7
Trace complete.
```

VDOMs in transparent mode

A VDOM in transparent mode is installed between the internal network and the router. In this mode, the VDOM does not make any changes to IP addresses and only applies security scanning to traffic. When a VDOM is added to a network in transparent mode, no network changes are required, except to provide the VDOM with a management IP address.

Each VDOM on a FortiGate can be configured for NAT/route mode or transparent mode, regardless of the operation mode of other VDOMs on the FortiGate. For more information about NAT/route mode, see ["VDOMs in NAT mode" on page 2949](#).

This chapter includes the following sections:

- [Transparent mode overview](#)
- [Using a VDOM in transparent mode](#)
- [VDOMs in transparent mode](#)

Transparent mode overview

In transparent mode, a VDOM becomes a layer-2 IP forwarding bridge. This means that Ethernet frames are forwarded based on destination MAC address, and no other routing is performed. All incoming traffic that is accepted by the firewall, is broadcast out on all interfaces.

In transparent mode the VDOM is a forwarding bridge, not a switch. A switch can develop a port table and associated MAC addresses, so that it can bridge two ports to deliver the traffic instead of broadcasting to all ports. In transparent mode, the VDOM does not follow this switch behavior, but instead is the forwarding bridge that broadcasts all packets out over all interfaces, subject to security policies.

Differences between NAT/route and transparent mode

The differences between NAT/route mode and transparent mode include:

Differences between NAT/route and transparent modes

Features	NAT/route mode	Transparent mode
Specific Management IP address required	No	Yes
Perform Network Address Translation (NAT)	Yes	Yes
Stateful packet inspection	Yes	Yes
Layer-2 forwarding	Yes	Yes
Layer-3 routing	Yes	No

Features	NAT/route mode	Transparent mode
Unicast Routing / Policy Based routing	Yes	No
DHCP server	Yes	No
IPsec VPN	Yes	Yes
PPTP/L2TP VPN	Yes	No
SSL VPN	Yes	No
Security features	Yes	Yes
VLAN support	Yes	Yes - limited to VLAN trunks.
Ping servers (dead gateway detection)	Yes	No

To provide administrative access to a FortiGate or VDOM in transparent mode, you must define a management IP address and a gateway. This step is not required in NAT/route mode where you can access the FortiGate through the assigned IP address of any interface where administrative access is permitted.

If you incorrectly set the transparent mode management IP address for your FortiGate, you will be unable to access your unit through the GUI. In this situation, you will need to connect to the FortiGate using the console cable and change the settings so you can access the unit. Alternately, if your unit has an LCD panel, you can change the operation mode and interface information through the LCD panel.

Operation mode differences in VDOMs

A VDOM, such as root, can have a maximum of 255 interfaces in Network Address Translation (NAT) mode or transparent mode. This includes VLANs, other virtual interfaces, and physical interfaces. To have more than a total of 255 interfaces configured, you need multiple VDOMs with multiple interfaces on each.

In transparent mode without VDOMs enabled, all interfaces on the FortiGate act as a bridge — all traffic coming in on one interface is sent back out on all the other interfaces. This effectively turns the FortiGate into a two interface unit no matter how many physical interfaces it has. When VDOMs are enabled, this allows you to determine how many interfaces to assign to a VDOM running in transparent mode. If there are reasons for assigning more than two interfaces based on your network topology, you're able to. However, the benefit of VDOMs in this case is that you have the functionality of transparent mode, but you can use interfaces for NAT/route traffic as well.

You can add more VDOMs to separate groups of VLAN subinterfaces. When using a FortiGate to serve multiple organizations, this configuration simplifies administration because you see only the security policies and settings for the VDOM you're configuring.

One essential application of VDOMs is to prevent problems caused when a FortiGate is connected to a layer-2 switch that has a global MAC table. FortiGate devices normally forward ARP requests to all interfaces, including VLAN subinterfaces. It is then possible for the switch to receive duplicate ARP packets on different VLANs. Some layer-2 switches reset when this happens. As ARP requests are only forwarded to interfaces in the same VDOM, you can solve this problem by creating a VDOM for each VLAN.

For more information about transparent mode, see the *Transparent Mode Handbook*.

Using a VDOM in transparent mode

The essential steps to configure a VDOM in transparent mode are:

- [Switching to transparent mode](#)
- [Adding VLAN subinterfaces](#)
- [Creating security policies](#)

You can also configure the security profiles that manage antivirus scanning, web filtering and spam filtering.

In transparent mode, you can access the GUI by connecting to an interface configured for administrative access and using HTTPS to access the management IP address. In the following examples, administrative access is enabled by default on the internal interface and the default management IP address is 10.11.0.1.

Switching to transparent mode

A VDOM is in NAT/route mode by default when it is created. You must switch it to transparent mode, and add a management IP address so you can access the VDOM from your management computer.



Before applying the change to transparent mode, ensure the VDOM has administrative access on the selected interface, and that the selected management IP address is reachable on your network.

Switching the VDOM to transparent mode can't be done through the GUI. It must be done through the CLI only.

To switch the VDOM to transparent mode - CLI:

```
config vdom
  edit <name>
    config system settings
      set opmode transparent
      set manageip 10.11.0.99 255.255.255.0
    end
  end
```

Adding VLAN subinterfaces

There are a few differences when adding VLANs in transparent mode compared to NAT/route mode.

In transparent mode, VLAN traffic is trunked across the VDOM. That means VLAN traffic can't be routed, changed, or inspected. For this reason when you assign a VLAN to a transparent mode VDOM, you will see the **Addressing Mode** section of the interface configuration disappear in from the GUI. It is because with no routing, inspection, or any activities able to be performed on VLAN traffic the VDOM simply re-broadcasts the VLAN traffic. This requires no addressing.

Also any routing related features such as dynamic routing or Virtual Router Redundancy Protocol (VRRP) are not available in transparent mode for any interfaces.

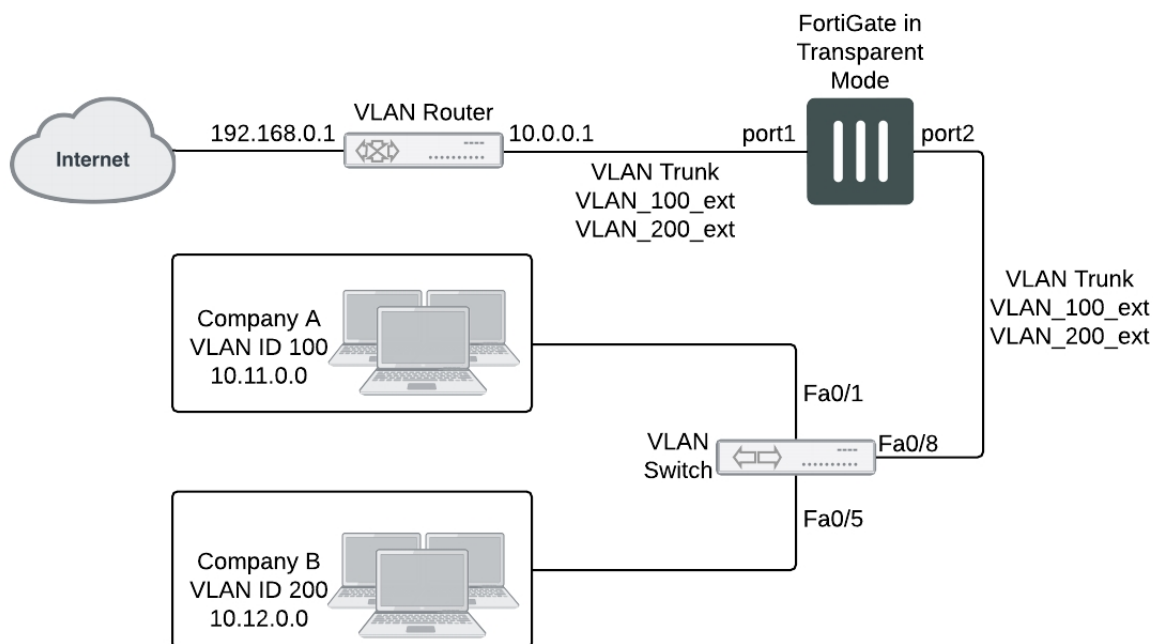
Creating security policies

Security policies permit communication between the FortiGate's network interfaces based on source and destination IP addresses. Typically you will also limit communication to desired times and services for additional security.

In transparent mode, the FortiGate performs antivirus and antispam scanning on each packet as it passes through the unit. You need security policies to permit packets to pass from the VLAN interface where they enter the unit to the VLAN interface where they exit the unit. If there are no security policies configured, no packets will be allowed to pass from one interface to another.

For more information, see the Firewall handbook.

Example configuration: VDOM in transparent mode



In this example, the FortiGate provides network protection to two organizations — Company A and Company B. Each company has different policies for incoming and outgoing traffic, requiring three different security policies and protection profiles.

VDOMs are not required for this configuration, but by using VDOMs the profiles and policies can be more easily managed on a per-VDOM basis either by one central administrator or separate administrators for each company. Also future expansion is simply a matter of adding additional VDOMs, whilst not disrupt the existing VDOMs.

For this example, firewalls are only included to deal with web traffic. This is to provide an example without making configuration unnecessarily complicated.

This example includes the following sections:

- [Network topology and assumptions](#)
- [General configuration steps](#)
- [Configuring common items](#)
- [Creating virtual domains](#)
- [Configuring the Company_A VDOM](#)
- [Configuring the Company_B VDOM](#)

- [Configuring the VLAN switch and router](#)
- [Testing the configuration](#)

Network topology and assumptions

Each organization's internal network consists of a different range of IP addresses:

- 10.11.0.0/255.255.0.0 for Company A.
- 10.12.0.0/255.255.0.0 for Company B.

For the procedures in this section, it is assumed that you have enabled VDOM configuration on your FortiGate. For more information, see [VDOMs overview](#).

The VDOM names are similar to the company names for easy recognition. The root VDOM can't be renamed and is not used in this example.

Interfaces used in this example are port1 and port2. Some FortiGate models may not have interfaces with these names. port1 is an external interface. port2 is an internal interface.

General configuration steps

The following steps summarize the configuration for this example. For best results, follow the procedures in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. [Configuring common items](#)
2. [Creating virtual domains](#)
3. [Configuring the Company_A VDOM](#)
4. [Configuring the Company_B VDOM](#)
5. [Configuring the VLAN switch and router](#)
6. [Testing the configuration](#)

Configuring common items

Both VDOMs require you configure security profiles. These will be configured the same way, but need to be configured in both VDOMs.

The relaxed profile allows users to surf websites they are not allowed to visit during normal business hours. Also a quota is in place to restrict users to one hour of access to these websites to ensure employees don't take long and unproductive lunches.

To create a strict web filtering profile - GUI:

1. Go to the proper VDOM, and select **Security Profiles > Web Filter**.
2. Select **Create New**.
3. Enter `strict` for the **Name**.
4. Expand FortiGuard Web Filtering, and select block for all Categories except Business Oriented, and Other.
5. Block all Classifications except Cached Content, and Image Search.
6. Ensure **FortiGuard Quota** for all Categories and Classifications is Disabled.
7. Select **OK**.

To create a strict web filtering profile - CLI:

```

config vdom
  edit <vdom_name>
    config webfilter profile
      edit strict
        config ftgd-wf
          set allow g07 g08 g21 g22 c01 c03
          set deny g01 g02 g03 g04 g05 g06 c02 c04 c05 c06 c07
        end
        set web-ftgd-err-log enable
      end
    end
  end
end

```

To create a relaxed web filtering profile - GUI:

1. Go to the proper VDOM, and select **Security Profiles > Web Filter**.
2. Select **Create New**.
3. Enter `relaxed` for the **Name**.
4. Expand FortiGuard Web Filtering, and select block for Potentially Security Violating Category, and Spam URL Classification.
5. Enable FortiGuard Quotas to allow 1 hour for all allowed Categories and Classifications.

Creating virtual domains

The FortiGate supports 10 virtual domains. Root is the default VDOM. It can't be deleted or renamed. The root VDOM is not used in this example. New VDOMs are created for Company A and Company B

To create the virtual domains - GUI:

1. With VDOMs enabled, select **Global > System > VDOM**.
2. Select **Create New**.
3. Enter `Company_A` for Name, and select **OK**.
4. Select **Create New**.
5. Enter `Company_B` for Name, and select **OK**.

To create the virtual domains - CLI:

```

config system vdom
  edit Company_A
  next
  edit Company_B
end

```

Configuring the Company_A VDOM

This section describes how to add VLAN subinterfaces and configure security policies for the Company_A VDOM.

This section includes the following topics:

- [Adding VLAN subinterfaces](#)
- [Creating the Lunch schedule](#)
- [Configuring Company_A firewall addresses](#)
- [Creating Company_A security policies](#)

Adding VLAN subinterfaces

You need to create a VLAN subinterface on the port2 interface and another one on the port1 interface, both with the same VLAN ID.

To add VLAN subinterfaces - GUI:

1. Go to **Global > Network > Interfaces**.
2. Select **Create New**.
3. Enter the following information and select **OK**:

Name	VLAN_100_int
Interface	port2
VLAN ID	100
Virtual Domain	Company_A

4. Select **Create New**.
5. Enter the following information and select **OK**:

Name	VLAN_100_ext
Interface	port1
VLAN ID	100
Virtual Domain	Company_A

To add the VLAN subinterfaces - CLI:

```
config system interface
  edit VLAN_100_int
    set interface port2
    set vlanid 100
    set vdom Company_A
  next
  edit VLAN_100_ext
    set interface port1
    set vlanid 100
    set vdom Company_A
end
```

Creating the Lunch schedule

Both organizations have the same lunch schedule, but only Company A has relaxed its security policy to allow employees more freedom in accessing the Internet during lunch. Lunch schedule will be Monday to Friday from 11:45 AM to 2:00 PM (14:00).

To create a recurring schedule for lunchtime - GUI:

1. In Company_A VDOM, go to **Policy & Objects > Schedules**.
2. Select **Create New**.
3. Enter **Lunch** as the name for the schedule.
4. Select **Mon, Tues, Wed, Thu, and Fri**.
5. Set the **Start** time as **11:45** and set the **Stop** time as **14:00**.
6. Select **OK**.

To create a recurring schedule for lunchtime - CLI:

```
config vdom
  edit Company_A
    config firewall schedule recurring
      edit Lunch
        set day monday tuesday wednesday thursday friday
        set start 11:45
        set end 14:00
      end
    end
```

Configuring Company_A firewall addresses

For Company A, its networks are all on the 10.11.0.0 network, so restricting addresses to that domain provides added security.

To configure Company_A firewall addresses - GUI:

1. In the Company_A VDOM, go to **Policy & Objects > Addresses**.
2. Select **Create New**.
3. Enter **CompanyA** in the **Address Name** field.
4. Type **10.11.0.0/255.255.0.0** in the **Subnet / IP Range** field.
5. Select **OK**.

To configure vdomA firewall addresses - CLI:

```
config firewall address
  edit CompanyA
    set type ipmask
    set subnet 10.11.0.0 255.255.0.0
  end
```

Creating Company_A security policies

A security policy can include varying levels of security feature protection. This example only deals with web filtering. The following security policies use the custom security `strict` and `relaxed` profiles configured earlier.

For these security policies, we assume that all protocols will be on their standard ports, such as port 80 for HTTP traffic. If the ports are changed, such as using port 8080 for HTTP traffic, you will have to create custom services for protocols with non-standard ports, and assign them different names.

The firewalls configured in this section are:

- internal to external — always allow all, security features - web filtering: strict
- internal to external — Lunch allow all, security features - web filtering: relaxed

Security policies allow packets to travel between the internal VLAN_100 interface to the external interface subject to the restrictions of the protection profile. Entering the policies in this order means the last one configured is at the top of the policy list, and will be checked first. This is important because the policies are arranged so if one does not apply the next is checked until the end of the list.

To configure Company_A security policies - GUI:

1. Go to **Policy & Objects > IPv4 Policy**.
2. Select **Create New**.
3. Enter the following information and select **OK**:

Name	CompanyA-lunch
Incoming Interface	VLAN_100_int
Outgoing Interface	VLAN_100_ext
Source Address	CompanyA
Destination Address	all
Schedule	Lunch
Service	all
Action	ACCEPT
Security Features	enable
Web Filtering	relaxed

This policy provides relaxed protection during lunch hours — going from strict down to scan for protocol options and web filtering. AntiVirus and Email Filtering remain at strict for security — relaxing them would not provide employees additional access to the Internet and it would make the company vulnerable.

1. Select **Create New**.
2. Enter the following information and select **OK**:

Name	CompanyA-strict
Incoming Interface	VLAN_100_int
Outgoing Interface	VLAN_100_ext
Source Address	CompanyA
Destination Address	all
Schedule	always

Service	all
Action	ACCEPT
Security Features	enable
Web Filtering	strict

This policy enforces strict scanning at all times, while allowing all traffic. It ensures company policies are met for network security.

4. Verify that the policy list arranged **By Sequence** to make sure the CompanyA-lunch policy is located above the CompanyA-strict policy. If necessary, rearrange the policies so that the appropriate policy is applied to outgoing traffic.

To configure Company_A security policies - CLI:

```
config vdom
  edit Company_A
    config firewall policy
      edit 1
        set name "CompanyA-lunch"
        set srcintf VLAN_100_int
        set dstintf VLAN_100_ext
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule Lunch
        set webfiltering relaxed
      next
      edit 2
        set name "CompanyA-strict"
        set srcintf VLAN_100_int
        set dstintf VLAN_100_ext
        set srcaddr all
        set dstaddr all
        set action accept
        set schedule always
        set webfiltering strict
      end
    end
  end
```

Configuring the Company_B VDOM

This section describes how to add VLAN subinterfaces and configure security policies for the Company B VDOM.

This section includes the following topics:

- [Adding VLAN subinterfaces](#)
- [Creating Company_B service groups](#)
- [Configuring Company_B firewall addresses](#)
- [Configuring Company_B security policies](#)

Adding VLAN subinterfaces

You need to create a VLAN subinterface on the internal interface and another one on the external interface, both with the same VLAN ID.

To add VLAN subinterfaces - GUI:

1. Go to **Network > Interfaces**.
2. Select **Create New**.
3. Enter the following information and select **OK**:

Name	VLAN_200_int
Interface	port2
VLAN ID	200
Virtual Domain	Company_B

4. Select **Create New**.
5. Enter the following information and select **OK**:

Name	VLAN_200_ext
Interface	port1
VLAN ID	200
Virtual Domain	Company_B

To add the VLAN subinterfaces - CLI:

```
config system interface
  edit VLAN_200_int
    set interface internal
    set vlanid 200
    set vdom Company_B
  next
  edit VLAN_200_ext
    set interface external
    set vlanid 200
    set vdom Company_B
end
```

Creating Company_B service groups

Company_B does not want its employees to use any online chat software except NetMeeting, which the company uses for net conferencing. To simplify the creation of a security policy for this purpose, you create a service group that contains all of the services you want to restrict. A security policy can manage only one service or one group.

To create a chat service group - GUI:

1. Go to **Policy & Objects > Services** and select **Create New > Service Group**.
2. Enter `Chat` in the **Group Name** field.
3. For each of IRC, AOL, SIP-MSNmessenger and TALK, select the service in the **Available Services** list and select the right arrow to add it to the **Members** list.

If a particular service does not appear in the **Available Services** list, see the list in **Policy & Objects > Services**. Some services don't appear by default unless edited.

4. Select **OK**.

To create a games and chat service group - CLI:

```
config firewall service group
edit Chat
set member IRC SIP-MSNmessenger AOL TALK
end
```

Configuring Company_B firewall addresses

Company B's network is all in the 10.12.0.0 network. Security can be improved by only allowing traffic from IP addresses on that network.

To configure Company_B firewall address - GUI:

1. In the Company_B VDOM, go to **Policy & Objects > Addresses**.
2. Select **Create New**.
3. Enter `new` in the **Address Name** field.
4. Type `10.12.0.0/255.255.0.0` in the **Subnet / IP Range** field.
5. Select **OK**.

To configure Company_B firewall addresses - CLI:

```
config vdom
edit Company_B
config firewall address
edit all
set type ipmask
set subnet 10.12.0.0 255.255.0.0
end
```

Configuring Company_B security policies

Security policies allow packets to travel between the internal and external VLAN_200 interfaces subject to the restrictions of the protection profile.

To configure Company_B security policies - GUI:

1. Go to **Policy & Objects > IPv4 Policy**.
2. Select **Create New**.
3. Enter the following information and select **OK**:

Name	CompanyB-deny-games-chat
Incoming Interface	VLAN_200_int
Outgoing Interface	VLAN_200_ext
Source Address	all
Destination Address	all
Schedule	BusinessDay
Service	games-chat
Action	DENY

This policy prevents the use of network games or chat programs (except NetMeeting) during business hours.

4. Enter the following information and select **OK**:

Name	CompanyB-lunch
Incoming Interface	VLAN_200_int
Outgoing Interface	VLAN_200_ext
Source Address	all
Destination Address	all
Schedule	Lunch
Service	HTTP, DNS
Action	ACCEPT
Security Features	enable
Web Filter	relaxed

This policy relaxes the web category filtering during lunch hour.

5. Select **Create New**.
6. Enter the following information and select **OK**:

Name	CompanyB-strict
Incoming Interface	VLAN_200_int
Outgoing Interface	VLAN_200_ext
Source Address	all

Destination Address	all
Schedule	BusinessDay
Service	HTTP, DNS
Action	ACCEPT
Security Profiles	enabled
Web Filter	strict

This policy provides rather strict web category filtering during business hours.

7. Select **Create New**.
8. Enter the following information and select **OK**:

Name	CompanyB-after-hours
Incoming Interface	VLAN_200_int
Outgoing Interface	VLAN_200_ext
Source Address	all
Destination Address	all
Schedule	always
Service	ANY
Action	ACCEPT
Security Profiles	enabled
Web Filter	relaxed

Because it is last in the list, this policy applies to the times and services not covered in preceding policies. This means that outside of regular business hours, the Relaxed protection profile applies to email and web browsing, and online chat and games are permitted. Company B needs this policy because its employees sometimes work overtime. The other companies in this example maintain fixed hours and don't want any after-hours Internet access.

To configure Company_B security policies - CLI:

```
config firewall policy
  edit 1
    set name "CompanyB-deny-games-chat"
    set srcintf VLAN_200_int
    set srcaddr all
    set dstintf VLAN_200_ext
    set dstaddr all
    set schedule BusinessDay
    set service Games
    set action deny
```



```
next
edit 2
    set name "CompanyB-lunch"
    set srcintf VLAN_200_int
    set srcaddr all
    set dstintf VLAN_200_ext
    set dstaddr all
    set action accept
    set schedule Lunch
    set service HTTP
    set profile_status enable
    set profile Relaxed
next
edit 3
    set name "CompanyB-strict"
    set srcintf VLAN_200_int
    set srcaddr all
    set dstintf VLAN_200_ext
    set dstaddr all
    set action accept
    set schedule BusinessDay
    set service HTTP
    set profile_status enable
    set profile BusinessOnly
next
edit 4
    set name "CompanyB-after-hours"
    set srcintf VLAN_200_int
    set srcaddr all
    set dstintf VLAN_200_ext
    set dstaddr all
    set action accept
    set schedule always
    set service ANY
    set profile_status enable
    set profile Relaxed
end
```

Configuring the VLAN switch and router

The Cisco switch is the first VLAN device internal passes through, and the Cisco router is the last device before the Internet or ISP.

This section includes the following topics:

- [Configuring the Cisco switch](#)
- [Configuring the Cisco router](#)

Configuring the Cisco switch

On the Cisco Catalyst 2900 ethernet switch, you need to define the VLANs 100, 200 and 300 in the VLAN database, and then add configuration files to define the VLAN subinterfaces and the 802.1Q trunk interface.

Add this file to Cisco VLAN switch:

```
!
interface FastEthernet0/1
switchport access vlan 100
```

```

!
interface FastEthernet0/5
switchport access vlan 300
!
interface FastEthernet0/6
switchport trunk encapsulation dot1q
switchport mode trunk
!

```

Switch 1 has the following configuration:

Port 0/1	VLAN ID 100
Port 0/3	VLAN ID 200
Port 0/6	802.1Q trunk

Configuring the Cisco router

The configuration for the Cisco router in this example is the same as in the basic example, except we add VLAN_300. Each of the three companies has its own subnet assigned to it.

The IP addresses assigned to each VLAN on the router are the gateway addresses for the VLANs. For example, devices on VLAN_100 would have their gateway set to 10.11.0.1/255.255.0.0.

```

!
interface FastEthernet0/0
switchport trunk encapsulation dot1q
switchport mode trunk
!
interface FastEthernet0/0.1
encapsulation dot1Q 100
ip address 10.11.0.1 255.255.0.0
!
interface FastEthernet0/0.3
encapsulation dot1Q 200
ip address 10.12.0.1 255.255.0.0
!

```

The router has the following configuration:

Port 0/0.1	VLAN ID 100
Port 0/0.3	VLAN ID 200
Port 0/0	802.1Q trunk

Testing the configuration

Use diagnostic commands, such as `tracert`, to test traffic routed through the network.

You should test traffic between the internal VLANs as well as from the internal VLANs to the Internet to ensure connectivity.

For additional troubleshooting, see [Troubleshooting VDOMs](#).

This section includes the following topics:

- [Testing traffic from VLAN_100 to the Internet](#)
- [Testing traffic from VLAN_100 to VLAN_200](#)

Testing traffic from VLAN_100 to the Internet

In this example, a route is traced from VLANs to a host on the Internet. The route target is `www.example.com`.

From a host on VLAN_100, access a command prompt and enter this command:

```
C:\>tracert www.example.com
Tracing route to www.example.com [208.77.188.166]
over a maximum of 30 hops:
  1 <10 ms <10 ms <10 ms 10.100.0.1
  ...
 14 172 ms 141 ms 140 ms 208.77.188.166
Trace complete.
```

The number of steps between the first and the last hop, as well as their IP addresses, will vary depending on your location and ISP. However, all successful tracerts to `www.example.com` will start and end with these lines.

Repeat the tracert for VLAN_200.

The tracert for each VLAN will include the gateway for that VLAN as the first step. Otherwise, the tracert should be the same for each VLAN.

Testing traffic from VLAN_100 to VLAN_200

In this example, a route is traced between two internal networks. The route target is a host on VLAN_200. The Windows traceroute command `tracert` is used.

From VLAN_100, access a Windows command prompt and enter this command:

```
C:\>tracert 10.12.0.2
Tracing route to 10.12.0.2 over a maximum of 30 hops:
  1 <10 ms <10 ms <10 ms 10.100.0.1
  2 <10 ms <10 ms <10 ms 10.12.0.2
Trace complete.
```

You can repeat this for different routes in the topology. In each case the IP addresses will be the gateway for the starting VLAN, and the end point at the ending VLAN.

Security profiles and VDOMs

A single VDOM can use all the security features that are available to a FortiGate that does not use VDOMs. However, you can also create global security profiles which are available for use by multiple VDOMs. Both types of profiles can be used together in your configuration.

This chapter includes the following sections:

- [VDOM-level security profiles](#)
- [Global security profiles](#)
- [FortiGuard licensing](#)

VDOM-level security profiles

A single VDOM can use all the security features that are available to a FortiGate that does not use VDOMs.

For an administrator who can only access a single VDOM, security profiles are found in the same part of the GUI and the CLI as for a single FortiGate that does not use VDOMs.

For the global administrator, you must use the dropdown menu at the top of the GUI to access a specific VDOM. Once this is done, go to **Security Profiles** to configure VDOM-level profiles. In the CLI, these profiles can be configured under the command `config vdom`.

Global security profiles

Security profiles can be configured globally for use by multiple VDOMs, to avoid creating identical profiles for each VDOM individually. Global profiles are available for the following security features:

- Antivirus
- Application control
- Data leak prevention
- Intrusion protection
- Web filtering

Some security profile features, such as URL filters, are not available for use in a global profile.

The name for any global profile must start with "g-" for identification. Global profiles are available as read-only for VDOM-level administrators and can only be edited or deleted from within the global settings. Each security feature has at least one default global profile.

Global profiles are configured by going to **Global > Security Profiles** in the GUI or under the following `config global` commands in the CLI:

- `antivirus profile`
- `application list`
- `dlp sensor`
- `ips sensor`
- `webfilter profile`

Inter-VDOM routing

Inter-VDOM routing changes this allows VDOMs to communicate internally without using additional physical interfaces, using VDOM links. VDOM links are virtual interfaces that connect VDOMs. A VDOM link contains a pair of interfaces with each one connected to a VDOM, and forming either end of the inter-VDOM connection.

This chapter contains the following sections:

- [Benefits of inter-VDOM routing](#)
- [Configuring VDOM links](#)
- [Inter-VDOM configurations](#)
- [Dynamic routing over inter-VDOM links](#)
- [HA virtual clusters and VDOM links](#)
- [Example configuration: inter-VDOM routing](#)

Benefits of inter-VDOM routing

Inter-VDOM routing has a number of advantages over independent VDOM routing. These benefits include:

- [Freed-up physical interfaces](#)
- [More speed than physical interfaces](#)
- [Continued support for secure firewall policies](#)
- [Configuration flexibility](#)

Freed-up physical interfaces

Tying up physical interfaces on the FortiGate presents a problem. With a limited number of interfaces available, configuration options for the old style of communication between VDOMs are very limited. VLANs can be an answer to this, but they have some limitations.

For example, the FortiGate-800 has 8 physical ethernet ports. If they are assigned 2 per VDOM (one each for external and internal traffic) there can only be 4 VDOMs at most configured, not the 10 VDOMs the license will allow. Adding even one additional interface per VDOM to be used to communicate between VDOMs leaves only 2 VDOMs for that configuration, since it would required 9 interfaces for 3 VDOMs. Even using one physical interface for both external traffic and inter-VDOM communication would severely lower the available bandwidth for external traffic on that interface.

With the introduction of inter-VDOM routing, traffic can travel between VDOMs internally, freeing up physical interfaces for external traffic. Using the above example we can use the 4 VDOM configuration and all the interfaces will have their full bandwidth.

More speed than physical interfaces

Internal interfaces are faster than physical interfaces. Their speed depends on the FortiGate CPU and its load. That means that an inter-VDOM link interface will be faster than a outbound physical interface connected to another inbound physical interface.

However, while one virtual interface with normal traffic would be considerably faster than on a physical interface, the more traffic and more internal interfaces you configure, the slower they will become until they are slower than

the physical interfaces. CPU load can come from other sources such as AV or content scanning. This produces the same effect—internal interfaces such as inter-VDOM links will be slower.

Continued support for secure firewall policies

VDOMs help to separate traffic based on your needs. This is an important step in satisfying regulations that require proof of secure data handling. This is especially important to health, law, accounting, and other businesses that handle sensitive data every day.

By keeping things separate, traffic has to leave the FortiGate and re-enter to change VDOMs. This forces traffic to go through the firewall when leaving and enter through another firewall, keeping traffic secure.

With inter-VDOM routing, the need for the physical interfaces is greatly reduced. However, firewall policies still need to be in place for traffic to pass through any interface, physical or virtual, and thus provide the same level of security both internally and externally. Configuration of firewall policies is the same for inter-VDOM links as for any other interface, and your data will continue to have the high level of security.

Configuration flexibility

A typical VDOM uses at least two interfaces, typically physical interfaces, one for internal and one for external traffic. Depending on the configuration, more interfaces may be required. This means that the maximum number of VDOMs configurable on a FortiGate using physical interfaces is the number of interfaces available divided by two. VLANs can increase the number by providing multiple virtual interfaces over a single physical interface, but VLANs have some limitations. Using physical interfaces for inter-VDOM communication therefore limits the number of possible configurations on your FortiGate.

To overcome this limitation, inter-VDOM links can be created within the FortiGate. Using virtual interfaces, inter-VDOM links free up the physical interfaces for external traffic. Using VDOM links on a FortiGate with 8 physical interfaces, you can have 4 VDOMs communicating with each other (meshed configuration) and continue to have 2 physical interfaces each for internal and external connections. This configuration would have required 20 physical interfaces without inter-VDOM routing. With inter-VDOM routing it only requires 8 physical interfaces, with the other 12 interfaces being internal VDOM links.

Inter-VDOM routing allows you to make use of standalone VDOMs, Management VDOMs, and Meshed VDOMs without being limited by the number of physical interfaces on your FortiGate. For more information about these types of VDOMs, see ["Inter-VDOM configurations" on page 2983](#).

Inter-VDOM configurations

By using fewer physical interfaces to inter-connect VDOMs, inter-VDOM links provide you with more configuration options.

None of these configurations use VLANs to reduce the number of physical interfaces. It is generally assumed that an internal or client network will have its own internal interface and an external interface to connect to its ISP and the Internet.


These inter-VDOM configurations can use any FortiGate model with possible limitations based on the number of physical interfaces. VLANs can be used to work around these limitations.

There are four different types of inter-VDOM configurations:

- [Standalone VDOM](#)
- [Independent VDOMs](#)

- Management VDOM
- Meshed VDOM

Standalone VDOM



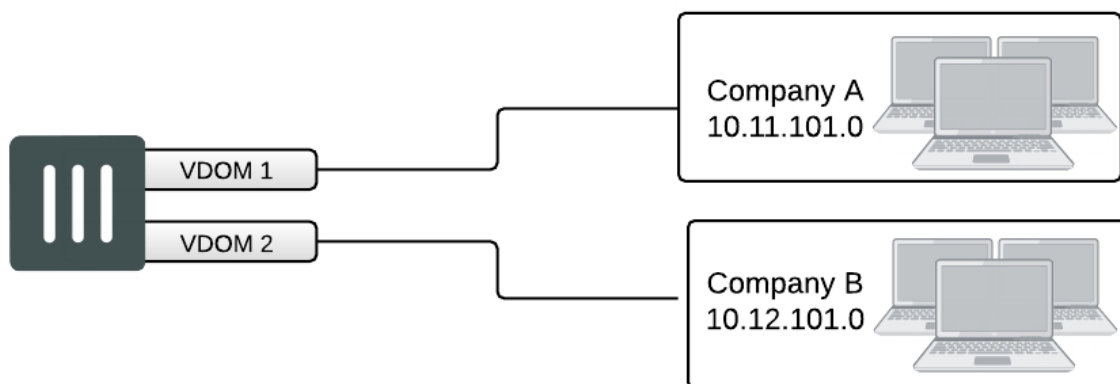
The standalone VDOM configuration uses a single VDOM on your FortiGate — the root VDOM that all FortiGate devices have by default. This is the VDOM configuration you're likely familiar with. It is the default configuration for FortiGate devices before you create additional VDOMs.

The configuration shown above has no VDOM inter-connections and requires no special configurations or settings.

The standalone VDOM configuration can be used for simple network configurations that only have one department or one company administering the connections, firewalls and other VDOM-dependent settings.

However, with this configuration, keeping client networks separate requires many interfaces, considerable firewall design and maintenance, and can quickly become time consuming and complex. Also, configuration errors for one client network can easily affect other client networks, causing unnecessary network downtime.

Independent VDOMs




The independent VDOMs configuration uses multiple VDOMs that are completely separate from each other. This is another common VDOM configuration.

This configuration has no communication between VDOMs and apart from initially setting up each VDOM, it requires no special configurations or settings. Any communication between VDOMs is treated as if communication is between separate physical devices.

The independent inter-VDOM configuration can be used where more than one department or one company is sharing the FortiGate. Each can administer the connections, firewalls and other VDOM-dependent settings for only its own VDOM. To each company or department, it appears as if it has its own FortiGate. This configuration reduces the amount of firewall configuration and maintenance required by dividing up the work.

However, this configuration lacks a management VDOM for VDOMs 1, 2, and 3. This management VDOM would enable an extra level of control for the FortiGate administrator, while still allowing each company or department to administer its own VDOM.

Management VDOM



In the management VDOM configuration, the root VDOM is the management VDOM. The other VDOMs are connected to the management VDOM with inter-VDOM links. There are no other inter-VDOM connections.

The inter-VDOM links connect the management VDOM to the other VDOMs. This does not require any physical interfaces, and the bandwidth of inter-VDOM links can be faster than physical interfaces, depending on the CPU workload.

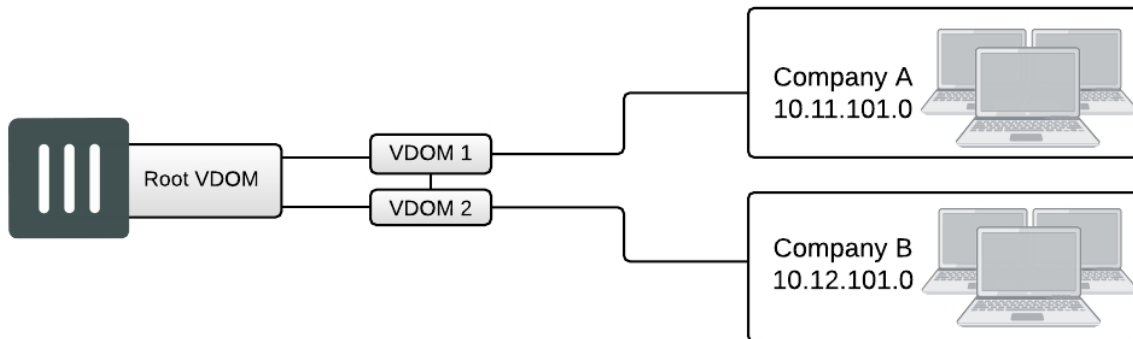
Only the management VDOM is connected to the Internet. The other VDOMs are connected to internal networks. All external traffic is routed through the management VDOM using inter-VDOM links and firewall policies between the management VDOM and each VDOM. This ensures the management VDOM has full control over access to the Internet, including what types of traffic are allowed in both directions. There is no communication directly between the non-root VDOMs. Security is greatly increased with only one point of entry and exit. Only the management VDOM needs to be fully managed to ensure network security in this case. Each client network can manage its own configuration without compromising security or bringing down another client network.

The management VDOM configuration is ideally suited for a service provider business. The service provider administers the management VDOM with the other VDOMs as customers. These customers don't require a dedicated IT person to manage their network. The service provider controls the traffic and can prevent the customers from using banned services and prevent Internet connections from initiating those same banned services. One example of a banned service might be Instant Messaging (IM) at a company concerned about intellectual property. Another example could be to limit bandwidth used by file-sharing applications without

banning that application completely. Firewall policies control the traffic between the customer VDOM and the management VDOM and can be customized for each customer.

The management VDOM configuration is limited in that the customer VDOMs have no inter-connections. In many situations this limitation is ideal because it maintains proper security. However, some configurations may require customers to communicate with each other, which would be easier if the customer VDOMs were inter-connected.

Meshed VDOM



The meshed VDOMs configuration, including partial and full mesh, has VDOMs inter-connected with other VDOMs. There is no special feature to accomplish this—they are just complex VDOM configurations.

Partial mesh means only some VDOMs are inter-connected. In a full mesh configuration, all VDOMs are inter-connected to all other VDOMs. This can be useful when you want to provide full access between VDOMs but handle traffic differently depending on which VDOM it originates from or is going to.

With full access between all VDOMs being possible, it is extra important to ensure proper security. You can achieve this level of security by establishing extensive firewall policies and ensuring secure account access for all administrators and users.

Meshed VDOM configurations can become complex very quickly, with full mesh VDOMs being the most complex. Ensure this is the proper solution for your situation before using this configuration. Generally, these configurations are seen as theoretical and are rarely deployed in the field.

Configuring VDOM links

Once VDOMs are configured on your FortiGate, configuring inter-VDOM routing and VDOM-links is very much like creating a VLAN interface. VDOM-links are managed through the GUI or CLI. In the GUI, VDOM link interfaces are managed in the network interface list.

This section includes the following topics:

- [Creating VDOM links](#)
- [IP addresses and inter-VDOM links](#)
- [Deleting VDOM links](#)
- [NAT to transparent VDOM links](#)

Creating VDOM links

VDOM links connect VDOMs together to allow traffic to pass between VDOMs as per firewall policies. Inter-VDOM links are virtual interfaces that are very similar to VPN tunnel interfaces except inter-VDOM links don't require IP addresses.

To create a VDOM link, you first create the point-to-point interface, and then bind the two interface objects associated with it to the virtual domains.

In creating the point-to-point interface, you also create two additional interface objects by default. They are called `vlink10` and `vlink11` - the interface name you chose with a 1 or a 0 to designate the two ends of the link.

Once the interface objects are bound, they are treated like normal FortiGate interfaces and need to be configured just like regular interfaces.

The assumptions for this example are as follows:

- Your FortiGate has VDOMs enabled and you have 2 VDOMs called `customer1` and `customer2` already configured. For more information on configuring VDOMs see [Configuring additional VDOMs](#).
- you're using a `super_admin` account.

To configure an inter-VDOM link - GUI:

1. Go to **Global > Network > Interfaces**.
2. Select **Create New > VDOM link**, enter the following information, and select **OK**.

Name	vlink1
	(The name can be up to 11 characters long. Valid characters are letters, numbers, "-", and "_". No spaces are allowed.)
Interface #0	
Virtual Domain	customer1
IP/Netmask	10.11.12.13/255.255.255.0
Administrative Access	HTTPS, SSL
Interface #1	
Virtual Domain	customer2
IP/Netmask	172.120.100.13/255.255.255.0
Administrative Access	HTTPS, SSL

To configure an inter-VDOM link - CLI:

```
config global
  config system vdom-link
    edit vlink1
  end
  config system interface
    edit vlink10
      set vdom customer1
```

```
next
edit vlink11
    set vdom customer2
end
```

Once you have created and bound the interface ends to VDOMs, configure the appropriate firewall policies and other settings that you require. To confirm the inter-VDOM link was created, find the VDOM link pair and use the expand arrow to view the two VDOM link interfaces. You can select edit to change any information.

IP addresses and inter-VDOM links

Besides being virtual interfaces, here is one main difference between inter-VDOM links and regular interfaces—default inter-VDOM links don't require IP addresses. IP addresses are not required by default because an inter-VDOM link is an internal connection that can be referred to by the interface name in firewall policies, and other system references. This introduces three possible situations with inter-VDOM links that are:

- **unnumbered** - an inter-VDOM link with no IP addresses for either end of the tunnel
- **half numbered** - an inter-VDOM link with one IP address for one end and none for the other end
- **full numbered** - an inter-VDOM link with two IP addresses, one for each end.

Not using an IP address in the configuration can speed up and simplify configuration for you. Also you will not use up all the IP addresses in your subnets if you have many inter-VDOM links.

Half or full numbered interfaces are required if you're doing NAT, either SNAT or DNAT as you need an IP number on both ends to translate between.

You can use unnumbered interfaces in static routing, by naming the interface and using 0.0.0.0 for the gateway. Running traceroute will not show the interface in the list of hops. However you can see the interface when you're sniffing packets, which is useful for troubleshooting.

Deleting VDOM links

When you delete the VDOM link, the two link objects associated with it will also be deleted. You can't delete the objects by themselves. The example uses a VDOM routing connection called "vlink1". Removing vlink1 will also remove its two link objects vlink10 and vlink11.



Before deleting the VDOM link, ensure all policies, firewalls, and other configurations that include the VDOM link are deleted, removed, or changed to no longer include the VDOM link.

To remove a VDOM link - GUI:

1. Go to **Global > Network > Interfaces**.
2. Select **Delete** for the VDOM link **vlink1**.

To remove a VDOM link - CLI:

```
config global
    config system vdom-link
        delete vlink1
    end
```

NAT to transparent VDOM links

Inter-VDOM links can be created between VDOMs in NAT mode and VDOMs in transparent mode, but it must be done through the CLI, as the VDOM link type must be changed from the default PPP to Ethernet for the two VDOMs to communicate. The below example assumes one vdom is in NAT mode and one is transparent mode.



An IP address must be assigned to the NAT VDOM interface, but no IP address should be assigned to the transparent VDOM interface.

To configure a NAT to transparent VDOM link - CLI:

```
config global
  config system vdom-link
    edit vlink1
      set type ethernet
    end
  config system interface
    edit vlink10
      set vdom (interface 1 name)
      set ip (interface 1 ip)
    next
    edit vlink11
      set vdom (interface 2 name)
    end
```

Ethernet-type is not recommended for standard NAT to NAT inter-VDOM links, as the default PPP-type link does not require the VDOM links to have addresses, while Ethernet-type does. VDOM link addresses are explained in [IP addresses and inter-VDOM links](#).

Dynamic routing over inter-VDOM links

BGP is supported over inter-VDOM links. Unless otherwise indicated, routing works as expected over inter-VDOM links.

If an inter-VDOM link has no assigned IP addresses to it, it may be difficult to use that interface in dynamic routing configurations. For example BGP requires an IP address to define any BGP router added to the network.

In OSPF, you can configure a router using a router ID and not its IP address. In fact, having no IP address avoids possible confusing between which value is the router ID and which is the IP address. However for that router to become adjacent with another OSPF router it will have to share the same subnet, which is technically impossible without an IP address. For this reason, while you can configure an OSPF router using an IP-less inter-VDOM link, it will likely be of limited value to you.

In RIP the metric used is hop count. If the inter-VDOM link can reach other nodes on the network, such as through a default route, then it may be possible to configure a RIP router on an inter-VDOM link. However, once again it may be of limited value due to limitations.

As stated earlier, BGP requires an IP address to define a router — an IP-less inter-VDOM link will not work with BGP.

In Multicast, you can configure an interface without using an IP address. However that interface will be unable to become an RP candidate. This limits the roles available to such an interface.

HA virtual clusters and VDOM links

FortiGate HA is implemented by configuring two or more FortiGate devices to operate as an HA cluster. To the network, the HA cluster appears to function as a single FortiGate, processing network traffic and providing normal security services such as firewall, VPN, IPS, virus scanning, web filtering, and spam filtering.

Virtual clustering extends HA features to provide failover protection and load balancing for a FortiGate operating with virtual domains. A virtual cluster consists of a cluster of two FortiGate devices operating with virtual domains. Traffic on different virtual domains can be load balanced between the cluster units.

With virtual clusters (vclusters) configured, inter-VDOM links must be entirely within one vcluster. You can't create links between vclusters, and you can't move a VDOM that is linked into another virtual cluster. If your FortiGate devices are operating in HA mode, with multiple vclusters when you create the vdom-link, the CLI command `config system vdom-link` includes an option to set which vcluster the link will be in.

What is virtual clustering?

Virtual clustering is an extension of the FGCP for FortiGate devices operating with multiple VDOMs enabled. Virtual clustering operates in active-passive mode to provide failover protection between two instances of a VDOM operating on two different cluster units. You can also operate virtual clustering in active-active mode to use HA load balancing to load balance sessions between cluster units. Alternatively, by distributing VDOM processing between the two cluster units you can also configure virtual clustering to provide load balancing by distributing sessions for different VDOMs to each cluster unit.

Virtual clustering and failover protection

Virtual clustering operates on a cluster of two (and only two) FortiGate devices with VDOMs enabled. Each VDOM creates a cluster between instances of the VDOMs on the two FortiGate devices in the virtual cluster. All traffic to and from the VDOM stays within the VDOM and is processed by the VDOM. One cluster unit is the primary unit for each VDOM and one cluster unit is the subordinate unit for each VDOM. The primary unit processes all traffic for the VDOM. The subordinate unit does not process traffic for the VDOM. If a cluster unit fails, all traffic fails over to the cluster unit that is still operating.

Virtual clustering and heartbeat interfaces

The HA heartbeat provides the same HA services in a virtual clustering configuration as in a standard HA configuration. One set of HA heartbeat interfaces provides HA heartbeat services for all of the VDOMs in the cluster. You don't have to add a heartbeat interface for each VDOM.

Virtual clustering and HA override

For a virtual cluster configuration, override is enabled by default for both virtual clusters when you:

- Enable VDOM portioning from the GUI by moving virtual domains to virtual cluster 2
- Enter `set vcluster2 enable` from the CLI `config system ha` command to enable virtual cluster 2.

Usually you would enable virtual cluster 2 and expect one cluster unit to be the primary unit for virtual cluster 1 and the other cluster unit to be the primary unit for virtual cluster 2. For this distribution to occur override must be enabled for both virtual clusters. Otherwise you will need to restart the cluster to force it to renegotiate.

Virtual clustering and load balancing or VDOM partitioning

There are two ways to configure load balancing for virtual clustering. The first is to set the HA mode to active-active. The second is to configure VDOM partitioning. For virtual clustering, setting the HA Mode to active-active has the same result as active-active HA for a cluster without virtual domains. The primary unit receives all sessions and load balances them among the cluster units according to the load balancing schedule. All cluster units process traffic for all virtual domains.

Note: If override is enabled the cluster may renegotiate too often. You can choose to disable override at any time. If you decide to disable override, for best results, you should disable it for both cluster units.


In a VDOM partitioning virtual clustering configuration, the HA mode is set to active-passive. Even though virtual clustering operates in active-passive mode you can configure a form of load balancing by using VDOM partitioning to distribute traffic between both cluster units. To configure VDOM partitioning you set one cluster unit as the primary unit for some virtual domains and you set the other cluster unit as the primary unit for other virtual domains. All traffic for a virtual domain is processed by the primary unit for that virtual domain. You can control the distribution of traffic between the cluster units by adjusting which cluster unit is the primary unit for each virtual domain.

For example, you could have 4 VDOMs, two of which have a high traffic volume and two of which have a low traffic volume. You can configure each cluster unit to be the primary unit for one of the high volume VDOMs and one of the low volume VDOMs. As a result each cluster unit will be processing traffic for a high volume VDOM and a low volume VDOM, resulting in an even distribution of traffic between the cluster units. You can adjust the distribution at any time. For example, if a low volume VDOM becomes a high volume VDOM you can move it from one cluster unit to another until the best balance is achieved. From the GUI you configure VDOM partitioning by setting the HA mode to active-passive and distributing virtual domains between Virtual Cluster 1 and Virtual Cluster 2. You can also configure different device priorities, port monitoring, and remote link failover, for Virtual Cluster 1 and Virtual Cluster 2.

From the CLI you configure VDOM partitioning by setting the HA mode to a-p. Then you configure device priority, port monitoring, and remote link failover and specify the VDOMs to include in virtual cluster 1. You do the same for virtual cluster 2 by entering the config secondary-vcluster command.

Failover protection does not change. If one cluster unit fails, all sessions are processed by the remaining cluster unit. No traffic interruption occurs for the virtual domains for which the still functioning cluster unit was the primary unit. Traffic may be interrupted temporarily for virtual domains for which the failed unit was the primary unit while processing fails over to the still functioning cluster unit. If the failed cluster unit restarts and rejoins the virtual cluster, VDOM partitioning load balancing is restored.

Example configuration: inter-VDOM routing



This example shows how to configure a FortiGate to use inter-VDOM routing.

This section contains the follow topics:

- [Network topology and assumptions](#)
- [Creating the VDOMs](#)
- [Configuring the physical interfaces](#)
- [Configuring the VDOM links](#)
- [Configuring the firewall and security profile settings](#)
- [Testing the configuration](#)

Network topology and assumptions

Two departments of a company, Accounting and Sales, are connected to one FortiGate 800 unit. To do its work, the Sales department receives a lot of email from advertising companies that would appear to be spam if the Accounting department received it. For this reason, each department has its own VDOM to keep firewall policies and other configurations separate. A management VDOM makes sense to ensure company policies are followed for traffic content.

The traffic between Accounting and Sales will be email and HTTPS only. It could use a VDOM link for a meshed configuration, but we will keep from getting too complex. With the configuration, inter-VDOM traffic will have a slightly longer path to follow than normal—from one department VDOM, through the management VDOM, and back to the other department VDOM. Since inter-VDOM links are faster than physical interfaces, this longer path should not be noticed.

Firewall policies will be in place. For added security, firewall policies will allow only valid office services such as email, web browsing, and FTP between either department and the Internet. Any additional services that are required can be added in the future.

The company uses a single ISP to connect to the Internet. The ISP uses DHCP to provide an IP address to the FortiGate. Both departments use the same ISP to reach the Internet.

Other assumptions for this example are as follows:

- Your FortiGate has interfaces labeled port1 through port4 and VDOMs are not enabled.
- you're using the super_admin account.
- You have the FortiClient application installed.
- you're familiar with configuring interfaces, firewalls, and other common features on your FortiGate.

General configuration steps

This example includes the following general steps. For best results, follow the steps in the order given. Also, note that if you perform any additional actions between procedures, your configuration may have different results.

1. [Creating the VDOMs](#)
2. [Configuring the physical interfaces](#)
3. [Configuring the VDOM links](#)
4. [Configuring the firewall and security profile settings](#)
5. [Testing the configuration](#)

Creating the VDOMs

This procedure enables VDOMs and creates the Sales and Accounting VDOMs.

To create the VDOMs - GUI:

1. Log in as the super_admin administrator.
2. Go to the **Dashboard** and locate the **System Information** widget. Enable **Virtual Domain**.
3. Log in again.
4. Go to **Global > System > VDOM**.
5. Select **Create New**, enter `Accounting` for the VDOM Name, and select **OK**.
6. Select **Create New**, enter `Sales` for the VDOM Name, and select **OK**.

To create the VDOMs - CLI:

```

config system global
    set vdom enable
end
config system vdom
    edit Accounting
    next
    edit Sales
    next
end

```

Configuring the physical interfaces

Next, the physical interfaces must be configured. This example uses three interfaces on the FortiGate - port2 (internal), port3 (dmz), and port1 (external). Port2 and port3 interfaces each have a department's network connected. port1 is for all traffic to or from the Internet and will use DHCP to configure its IP address, which is common with many ISPs.

To configure the physical interfaces - GUI:

1. Go to **Global > Network > Interfaces**.
2. Select **Edit** for the port2 interface, enter the following information, and select **OK**.

Alias	AccountingLocal
Virtual Domain	Accounting
Addressing mode	Manual
IP/Netmask	172.100.1.1/255.255.0.0
Administrative Access	HTTPS, PING, SSH
Description	This is the accounting department internal interface.

3. Select **Edit** for the port3 interface, enter the following information, and select **OK**.

Alias	SalesLocal
Virtual Domain	Sales
Addressing mode	Manual

IP/Netmask	192.168.1.1/255.255.0.0
Administrative Access	HTTPS, PING, SSH
Description	This is the sales department internal interface.

4. Select **Edit** for the port1 interface, enter the following information, and select **OK**.

Alias	ManagementExternal
Virtual Domain	root
Addressing Mode	DHCP
Distance	5
Retrieve default gateway from server	Enable
Override internal DNS	Enable
Administrative Access	HTTPS, SSH, SNMP
Description	This is the accounting department internal interface.



When the mode is set to DHCP or PPoE on an interface you can set the distance field. This is the administrative distance for any routes learned through the gateway for this interface. The gateway is added to the static route table with these values. A lower distance indicates a preferred route.

To configure the physical interfaces - CLI:

```
config global
config system interface
edit port2
set alias AccountingLocal
set vdom Accounting
set mode static
set ip 172.100.1.1 255.255.0.0
set allowaccess https ping ssh
set description "The accounting dept internal interface"
next
edit port3
set alias SalesLocal
set vdom Sales
set mode static
set ip 192.168.1.1 255.255.0.0
set allowaccess https ping ssh
set description "The sales dept. internal interface"
next
edit port1
set alias ManagementExternal
set vdom root
set mode DHCP
```

```

set distance 5
set gwdetect enable
set dns-server-override enable
set allowaccess https ssh snmp
set description "The system-wide management interface."
end

```

Configuring the VDOM links

To complete the connection between each VDOM and the management VDOM, you need to add the two VDOM links; one pair is the Accounting - management link and the other is for Sales - management link.

When configuring inter-VDOM links, you don't have to assign IP addresses to the links unless you're using advanced features such as dynamic routing that require them or want to add them for clarity. Not assigning IP addresses results in faster configuration, and more available IP addresses on your networks.

To configure the Accounting and management VDOM link - GUI:

1. Go to **Global > Network > Interfaces**.
2. Select the expand arrow to select **Create New > VDOM link**.
3. Enter the following information, and select **OK**.

Name	AccountVlnk
Interface #0	
Virtual Domain	Accounting
IP/Netmask	0.0.0.0/0.0.0.0
Administrative Access	HTTPS, PING, SSH
Description	The Accounting VDOM side of the link.
Interface #1	
Virtual Domain	root
IP/Netmask	0.0.0.0/0.0.0.0
Administrative Access	HTTPS, PING, SSH
Description	The Management VDOM side of the link.

To configure the Accounting and management VDOM link - CLI:

```

config global
  config system vdom-link
    edit AccountVlnk
    next
  end
  config system interface
    edit AccountVlnk0
      set vdom Accounting

```

```

        set ip 0.0.0.0 0.0.0.0
        set allowaccess https ping ssh
        set description "Accounting side of the VDOM link"
    next
    edit AccountVlnk1
        set vdom root
        set ip 0.0.0.0 0.0.0.0
        set allowaccess https ping ssh
        set description "Management side of the VDOM link"
    end

```

To configure the Sales and management VDOM link - GUI:

1. Go to **Global > Network > Interfaces**.
2. Select the expand arrow and select **Create New > VDOM link**.
3. Enter the following information, and select **OK**.

Name	SalesVlnk
Interface #0	
Virtual Domain	Sales
IP/Netmask	0.0.0.0/0.0.0.0
Administrative Access	HTTPS, PING, SSH
Description	The Sales VDOM side of the link.
Interface #1	
Virtual Domain	root
IP/Netmask	0.0.0.0/0.0.0.0
Administrative Access	HTTPS, PING, SSH
Description	The Management VDOM side of the link.

To configure the Sales and management VDOM link - CLI:

```

config global
    config system vdom-link
        edit SalesVlnk
    end
config system interface
    edit SalesVlnk0
        set vdom Accounting
        set ip 0.0.0.0 0.0.0.0
        set allowaccess https ping ssh
        set description "Sales side of the VDOM link"
    next
    edit SalesVlnk1
        set vdom root

```

```
set ip 0.0.0.0 0.0.0.0
set allowaccess https ping ssh
set description "Management side of the VDOM link"
end
end
```

Configuring the firewall and security profile settings

With the VDOMs, physical interfaces, and VDOM links configured the firewall must now be configured to allow the proper traffic. Firewalls are configured per-VDOM, and firewall objects must be created for each VDOM separately.

For this example, the firewall group of services allowed between the internal networks and the Internet are the basic services for web browsing, file transfer, and email. These include: HTTP, HTTPS, SSL, FTP, DNS, NTP, POP3, and SMTP.

The only services allowed between Sales and Accounting are secure web browsing (HTTPS) and email (POP3 and SMTP).



The limited number of services ensures security between departments. The list of services can be expanded in the future if needed.

Security profile settings will block all non-essential business websites while logging all web traffic, scan and file filter all web and email protocols, and block game and peer-to-peer applications using application control.

For added security, FortiClient is required on internal computers with AntiVirus scanning configured. This is enforced by **Endpoint NAC** in firewall policies.

Using firewall addresses makes the firewall policies easier to read. Also if any changes need to be made in the future, you can simply update the addresses without changing the firewall policies. The addresses required are:

- AccountingLocal - all traffic from the internal accounting network
- AccountingVlnk - all traffic from the VDOM link between accounting and management VDOMs
- SalesLocal - all traffic from the internal sales network
- SalesVlnk - all traffic from the VDOM link between sales and management VDOM.

The Accounting VDOM requires AccountingLocal, AccountingVlnk, and SalesLocal. The Sales VDOM requires SalesLocal, SalesVlnk, and AccountingLocal.

This section includes the following topics:

- [Configuring firewall service groups](#)
- [Configuring security profile settings for the Accounting VDOM](#)
- [Configuring firewall settings for the Accounting VDOM](#)
- [Configuring security profile settings for the Sales VDOM](#)
- [Configuring firewall settings for the Sales VDOM](#)
- [Configuring firewall settings between the Accounting and Sales VDOMs](#)

Configuring firewall service groups

Service groups are an easy way to manage multiple services, especially if the same services are used on different networks.

The two service groups used here are intended for normal office traffic to the Internet, and for restricted traffic between departments. In both cases network traffic will be limited to the services listed to prevent any potential security risks or bandwidth-robbing applications.

These service groups can be changed as needed to either include additional valid services that are being used on the network, or to exclude services that are not required. Also, custom services can be created as needed for applications that are not listed.

To configure two firewall service groups - GUI:

1. Open the **Accounting** VDOM.
2. Go to **Policy & Objects > Services** and select **Create New > Service Group**.
3. Select **Create New**, enter the following information, and select **OK**.

Group Name	OfficeServices
Members	HTTP, HTTPS, SSL, FTP, DNS, NTP, POP3, PING, SMTP

4. Select **Create New**, enter the following information, and select **OK**.

Group Name	AccountingSalesServices
Members	HTTPS, POP3, PING, SMTP

To configure two firewall service groups - CLI:

```
config vdom
  edit Accounting
    config firewall service group
      edit OfficeServices
        set member HTTP HTTPS SSH FTP DNS NTP POP3 PING SMTP
      next
      edit AccountingSalesServices
        set member HTTPS POP3 PING SMTP
      end
    end
  end
```

Configuring security profile settings for the Accounting VDOM

Security profile settings include web filtering, antivirus, application control, and other features. This example just uses those three features to ensure that:

- the business environment is free from viruses
- employees don't access inappropriate websites
- employees don't use games or peer-to-peer applications at work

To configure web filtering for the Accounting VDOM - GUI:

1. Open the **Accounting** VDOM.
2. Go to **Security Profiles > Web Filter**.
3. Select **Create New**.
4. Enter `webStrict` for the **Name**.

5. Select the arrow to expand the **FortiGuard Web Filtering** section.
6. Block all **Categories** except Business Oriented, Other, and Unrated.
7. Block all **Classifications** except Image Search.
8. Log all **Categories** and **Classifications**.
9. Select **OK**.

To configure AntiVirus for the Accounting VDOM - GUI:

1. Open the **Accounting** VDOM.
2. Go to **Security Profiles > AntiVirus**.
3. Select **Create New**.
4. Enter `avStrict` for the **Name**.
5. Set **Detect Viruses** to **Block** and enable all **Inspected Protocols**.
6. Select **OK**.

To configure application control for the Accounting VDOM - GUI:

1. Open the **Accounting** VDOM.
2. Go to **Security Profiles > Application Control**.
3. Select **Create New** (+ button at top right of page).
4. Enter `appStrict` for **Name** and select **OK**.
5. Select **Create New**.
6. In **Filters**, set **Category** to **game**.
7. In **Applications/Settings**, enter the following, and select **OK**.

Action	Block
Packet Logging	Enable

8. Select **Create New**.
9. In **Filters**, set **Category** to **p2p**.
10. In **Applications/Settings**, enter the following, and select **OK**.

Action	Block
Packet Logging	Enable

11. Select **Apply**.

To configure application control for the Accounting VDOM - CLI:

```
config vdom
  edit Accounting
    config application list
      edit appStrict
        config entries
          edit 1
            set category 2
          next
        
```

```

        edit 2
        set category 8
        end
    end
end

```

Configuring firewall settings for the Accounting VDOM

This configuration includes two firewall addresses and two firewall policies for the Accounting VDOM - one for the internal network, and one for the VDOM link with the management VDOM (root).

For added security, all traffic allowed will be scanned. Only valid office traffic will be allowed using the service group `OfficeServices`.

Note the spelling of `AccountVlnk` which is due to the eleven character limit on VDOM link names.

To configure firewall addresses - GUI:

1. Open the **Accounting** VDOM.
2. Select **Policy & Objects > Addresses**.
3. Select **Create New**, enter the following information, and select **OK**.

Address Name	AccountingLocal
Type	Subnet/ IP Range
Subnet / IP Range	172.100.0.0
Interface	port1

4. Select **Create New**, enter the following information, and select **OK**.

Address Name	AccountManagement
Type	Subnet/ IP Range
Subnet / IP Range	10.0.1.0
Interface	AccountVlnk

To configure firewall addresses - CLI:

```

config vdom
  edit Accounting
    config firewall address
      edit AccountingLocal
        set type iprange
        set subnet 172.100.0.0
        set associated-interface port1
      next
      edit AccountManagement
        set type iprange
        set subnet 10.0.1.0
        set associated-interface AccountVlnk
      end
    end
  end

```

end

To configure the firewall policies from AccountingLocal to the Internet - GUI:

1. Open the **Accounting** VDOM.
2. Go to **Policy & Objects > IPv4 Policy**.
3. Select **Create New**, enter the following information, and then select **OK**.

Name	Accounting-Local-to-Management
Incoming Interface	port2
Outgoing Interface	AccountVlnk
Source Address	AccountingLocal
Destination Address	AccountManagement
Schedule	always
Service	OfficeServices
Action	ACCEPT
Enable NAT	enable
Security Features	enabled
Web Filtering	webStrict
AntiVirus Filtering	avStrict
Application Control	appStrict

4. Open the **root** VDOM.
5. Go to **Policy & Objects > IPv4 Policy**.
6. Select **Create New**, enter the following information, and then select **OK**.

Name	Accounting-VDOM-to-Internet
Incoming Interface	AccountVlnk
Outgoing Interface	ManagementExternal
Source Address	AccountManagement
Destination Address	all
Schedule	always
Service	OfficeServices
Action	ACCEPT

Enable NAT	enable
Security Features	enabled
Web Filtering	webStrict
AntiVirus Filtering	avStrict
Application Control	appStrict

To configure the firewall policies from AccountingLocal to Internet - CLI:

```

config vdom
  edit Accounting
    config firewall policy
      edit 1
        set name "Accounting-Local-to-Management"
        set srcintf port2
        set dstintf AccountVlnk
        set srcaddr AccountingLocal
        set dstaddr AccountManagement
        set action accept
        set schedule always
        set service OfficeServices
        set nat enable
        set av-profile avStrict
        set webfilter-profile webStrict
        set application-list appStrict
      end
    end
  end
config vdom
  edit root
    config firewall policy
      edit 2
        set name "Accounting-VDOM-to-Internet"
        set srcintf AccountVlnk
        set dstintf port1
        set srcaddr AccountManagement
        set dstaddr all
        set action accept
        set schedule always
        set service OfficeServices
        set nat enable
        set av-profile scan
        set webfilter-profile scan
        set application-list AppControlList
      end
    end
  end

```

To configure the firewall policies from Internet to AccountingLocal - GUI:

1. Open the **root** VDOM.
2. Go to **Policy & Objects > IPv4 Policy**.
3. Select **Create New**, enter the following information, and select **OK**.

Name	Internet-access-to-Accounting-VDOM
Incoming Interface	port1
Outgoing Interface	AccountVlnk
Source Address	all
Destination Address	AccountManagement
Schedule	always
Service	OfficeServices
Action	ACCEPT
Enable NAT	enable
Security Features	enabled
Web Filtering	webStrict
AntiVirus Filtering	avStrict
Application Control	appStrict

4. Open the **Accounting** VDOM.
5. Go to **Policy & Objects > IPv4 Policy**.
6. Select **Create New**, enter the following information, and select **OK**.

Name	Management-access-to-Accounting-local
Incoming Interface	AccountVlnk
Outgoing Interface	port2
Source Address	AccountManagement
Destination Address	AccountingLocal
Schedule	always
Service	OfficeServices
Action	ACCEPT
Enable NAT	enable
Security Features	enabled
Web Filtering	webStrict
AntiVirus Filtering	avStrict
Application Control	appStrict

To configure the firewall policies from Internet to AccountingLocal - CLI:

```
config vdom
  edit root
    config firewall policy
      edit 3
        set name "Internet-access-to-Accounting-VDOM"
        set srcintf port1
        set dstintf AccountVlnk
        set srcaddr all
        set dstaddr AccountManagement
        set action accept
        set schedule always
        set service OfficeServices
        set nat enable
        set av-profile avStrict
        set webfilter-profile webStrict
        set application-list appstrict
      end
    end
  config vdom
    edit Accounting
      config firewall policy
        edit 4
          set name "Management-access-to-Accounting-local"
          set srcintf AccountVlnk
          set dstintf port2
          set srcaddr AccountManagement
          set dstaddr AccountingLocal
          set action accept
          set schedule always
          set service OfficeServices
          set nat enable
          set av-profile avStrict
          set webfilter-profile webStrict
          set application-list appstrict
        end
      end
    end
  end
```

Configuring security profile settings for the Sales VDOM

Security profile settings include web filtering, antivirus, application control, and other features. This example just uses those three features to ensure that

- the business environment is free from viruses
- employees don't surf grossly inappropriate websites, and
- employees don't use games or peer-to-peer applications at work.

Note that Sales web traffic is different from Accounting, and web filtering is different to account for this.

To configure web filtering for the Sales VDOM - GUI:

1. Open the **Sales** VDOM.
2. Go to **Security Profiles > Web Filter**.
3. Select **Create New**.
4. Enter `webStrict` for the **Name**.

5. In **FortiGuard Categories**, select all of the categories except **Bandwidth Consuming**, **General Interest - Business** and **Unrated**.
6. In **Change Action for Selected Categories** select **Block**.
7. Select **Apply**.

To configure web filtering for the Sales VDOM - CLI:

```
config vdom
  edit Sales
    config webfilter profile
      edit webStrict
        config ftgd-wf
          set allow g07 g08 g21 g22 c01 c03
          set deny g01 g02 g03 g04 g05 g06 c02 c04 c05 c06 c07
        end
        set web-ftgd-err-log enable
      end
    end
  end
```

To configure AntiVirus for the Sales VDOM - GUI:

1. Open the **Sales** VDOM.
2. Go to **Security Profiles > AntiVirus**.
3. Select **Create New**.
4. Enter `avStrict` for the **Name**.
5. Set **Detect Viruses** to **Block** and enable all **Inspected Protocols**.
6. Select **Apply**.

To configure AntiVirus for the Sales VDOM - CLI:

```
config vdom
  edit Sales
    config antivirus profile
      edit "avStrict"
        config http
          set options scan file-filter
        end
        config ftp
          set options scan file-filter
        end
        config imap
          set options scan file-filter
        end
        config pop3
          set options scan file-filter
        end
        config smtp
          set options scan file-filter
        end
        config nntp
          set options scan file-filter
        end
        config im
          set options scan file-filter
        end
      end
    end
  end
```

```

        end
        set filepattable 1
        set av-virus-log enable
        set av-block-log enable
    end
end

```

To configure application control for the Sales VDOM - GUI:

1. Open the **Accounting** VDOM.
2. Go to **Security Profiles > Application Control**.
3. Select **Create New** (+ button at top right of page).
4. Enter `appStrict` for **Name** and select **OK**.
5. Select **Create New**.
6. In **Filters**, set **Category** to **game**.
7. In **Applications/Settings**, enter the following, and select **OK**.

Action	Block
Packet Logging	Enable

8. Select **Create New**.
9. In **Filters**, set **Category** to **p2p**.
10. In **Applications/Settings**, enter the following, and select **OK**.

Action	Block
Packet Logging	Enable

11. Select **Apply**.

To configure application control for the Sales VDOM - CLI:

```

config vdom
  edit Sales
    config application list
      edit "appStrict"
        config entries
          edit 1
            set category 2
          next
          edit 2
            set category 8
          end
        end
      end
    end
  end
end

```

Configuring firewall settings for the Sales VDOM

Like the Accounting firewall settings, this configuration includes two firewall addresses and two firewall policies for the sales VDOM: one for the internal network, and one for the VDOM link with the management VDOM.

When entering the CLI commands, the number of the firewall policies must be high enough to be a new policy. Depending on the number of firewall policies on your FortiGate, this may require starting at a higher number than the 6 required for the default configuration. This number is added automatically when you configure firewall policies using the web manager interface.

The FortiClient application must be used on Sales network computers to ensure additional protection for the sensitive information and for protection against spam.

To configure firewall addresses - GUI:

1. Open the **Sales** VDOM.
2. Go to **Policy & Objects > Addresses**.
3. Select **Create New**, enter the following information, and select **OK**.

Address Name	SalesLocal
Type	Subnet / IP Range
Subnet / IP Range	172.100.0.0
Interface	port3

4. Go to **Policy & Objects > Addresses**.
5. Select **Create New**, enter the following information, and select **OK**.

Address Name	SalesManagement
Type	Subnet / IP Range
Subnet / IP Range	10.0.1.0
Interface	SalesVlnk

To configure the firewall addresses - CLI:

```
config vdom
  edit Sales
    config firewall address
      edit SalesLocal
        set type iprange
        set subnet 172.100.0.0
        set associated-interface port2
      next
      edit SalesManagement
        set type iprange
        set subnet 10.0.1.0
        set associated-interface SalesVlnk
      end
    end
  end
```

To configure the firewall policies from SalesLocal to the Internet - GUI:

1. Open the **Sales** VDOM.
2. Go to **Policy & Objects > IPv4 Policy**.

3. Select **Create New**, enter the following information, and select **OK**.

Name	Sales-local-to-Management
Incoming Interface	port3
Outgoing Interface	SalesVlnk
Source Address	SalesLocal
Destination Address	SalesManagement
Schedule	always
Service	OfficeServices
Action	ACCEPT
Log Allowed Traffic	enabled

4. Open the **root** VDOM.
 5. Go to **Policy & Objects > IPv4 Policy**.
 6. Select **Create New**, enter the following information, and select **OK**.

Name	Sales-VDOM-to-Internet
Incoming Interface	SalesVlnk
Outgoing Interface	ManagementExternal
Source Address	SalesManagement
Destination Address	all
Schedule	always
Service	OfficeServices
Action	ACCEPT
Log Allowed Traffic	enabled

To configure the firewall policies from SalesLocal to the Internet - CLI:

```
config vdom
  edit root
    config firewall policy
      edit 6
        set name "Sales-local-to-Management"
        set srcintf port2
        set srcaddr SalesLocal
        set dstintf SalesVlnk
        set dstaddr SalesManagement
        set schedule always
        set service OfficeServices
        set action accept
```

```

        set logtraffic enable
    end
end
config vdom
    edit Sales
        config firewall policy
            edit 7
                set name "Sales-VDOM-to-Internet"
                set srcintf SalesVlnk
                set srcaddr SalesManagement
                set dstintf external
                set dstaddr all
                set schedule always
                set service OfficeServices
                set action accept
                set logtraffic enable
            end
        end
    end
end

```

To configure the firewall policies from the Internet to SalesLocal - GUI:

1. Open the **root** VDOM.
2. Go to **Policy & Objects > IPv4 Policy**.
3. Select **Create New**, enter the following information, and select **OK**.

Name	Internet-access-to-Sales-VDOM
Incoming Interface	ManagementExternal
Outgoing Interface	SalesVlnk
Source Address	all
Destination Address	SalesManagement
Schedule	always
Service	OfficeServices
Action	ACCEPT
Protection Profile	scan
Log Allowed Traffic	enabled

4. Open the **Sales** VDOM.
5. Go to **Policy & Objects > IPv4 Policy**.
6. Select **Create New**, enter the following information, and select **OK**.

Name	Management-access-to-Sales-local
Incoming Interface	SalesVlnk

Outgoing Interface	port2
Source Address	SalesManagement
Destination Address	SalesLocal
Schedule	always
Service	OfficeServices
Action	ACCEPT
Log Allowed Traffic	enabled

To configure the firewall policies from the Internet to SalesLocal - CLI:

```

config vdom
  edit root
    config firewall policy
      edit 8
        set name "Internet-access-to-Sales-VDOM"
        set srcintf external
        set srcaddr all
        set dstintf SalesVlnk
        set dstaddr SalesManagement
        set schedule always
        set service OfficeServices
        set action accept
        set logtraffic enable
      end
    end
  config vdom
    edit Sales
      config firewall policy
        edit 9
          set name "Management-access-to-Sales-local"
          set srcintf SalesVlnk
          set srcaddr SalesManagement
          set dstintf port2
          set dstaddr SalesLocal
          set schedule always
          set service OfficeServices
          set action accept
          set logtraffic enable
        end
      end
    end
  end

```

Configuring firewall settings between the Accounting and Sales VDOMs

Firewall policies are required for any communication between each internal network and the Internet. Policies are also required for the two internal networks to communicate with each other through the management VDOM.

The more limited AccountingSalesServices group of services will be used between Sales and Accounting to ensure the traffic is necessary business traffic only. These policies will result in a partially meshed VDOM

configuration. The FortiClient application must be used to ensure additional protection for the sensitive accounting information.

Two firewall policies are required to allow traffic in both directions between Sales and Accounting.

To configure the firewall policy between Sales and Accounting on the management VDOM - GUI:

1. Open the **root** VDOM.
2. Go to **Policy & Objects > IPv4 Policy**.
3. Select **Create New**, enter the following information, and select **OK**.

Name	Sales-VDOM-to-Accounting-VDOM
Incoming Interface	SalesVlnk
Outgoing Interface	AccountVlnk
Source Address	SalesManagement
Destination Address	AccountingManagement
Schedule	always
Service	AccountingSalesServices
Action	ACCEPT
Protection Profile	scan
Log Allowed Traffic	enabled

4. Go to **Policy & Objects > IPv4 Policy**.
5. Select **Create New**, enter the following information, and select **OK**.

Name	Accounting-VDOM-to-Sales-VDOM
Incoming Interface	AccountVlnk
Outgoing Interface	SalesVlnk
Source Address	AccountingManagement
Destination Address	SalesManagement
Schedule	always
Service	AccountingSalesServices
Action	ACCEPT
Log Allowed Traffic	enabled

To configure the firewall policy between Sales and Accounting on the management VDOM - CLI:

```
config vdom
```

```
edit root
  config system firewall policy
    edit 9
      set name "Sales-VDOM-to-Accounting-VDOM"
      set srcintf SalesVlnk
      set srcaddr SalesManagement
      set dstintf AccountVlnk
      set dstaddr AccountManagement
      set schedule always
      set service AccountingSalesServices
      set action accept
      set logtraffic enable
    next
    edit 10
      set name "Accounting-VDOM-to-Sales-VDOM"
      set srcintf AccountVlnk
      set srcaddr AccountManagement
      set dstintf SalesVlnk
      set dstaddr SalesManagement
      set schedule always
      set service AccountingSalesServices
      set action accept
      set logtraffic enable
    end
  end
```

Testing the configuration

Once the inter-VDOM routing has been configured, tests must be conducted to confirm proper operation. If there are any problems, use the troubleshooting tips to resolve them.

This section includes the following topics:

- [Testing connectivity](#)
- [Troubleshooting Tips](#)

Testing connectivity

Testing connectivity ensures that physical networking connections as well as FortiGate interface configurations, including firewall policies, are properly configured.

The easiest way to test connectivity is to use the `ping` and `tracert` commands to confirm the connectivity of different routes on the network. Include testing:

- from AccountingLocal to Internet
- from Internet to AccountingLocal
- from SalesLocal to Internet
- from Internet to SalesLocal
- from AccountingLocal to SalesLocal.

When using the commands on a Windows computer, go to a command line prompt and enter either `ping <IP address>` or `tracert <IP address>`.

When using the commands on a FortiGate, go to the CLI and enter either `exec ping <IP address>` or `exec traceroute <IP address>`.

Troubleshooting Tips

When there are problems with connectivity, the following troubleshooting tips will help resolve the issues.

- If a multiple hop test, such as traceroute, is not successful then reduce it to a single hop to simplify the test. Test each link of the path to see which hop is down. If all hops are up, check the FortiGate policies to ensure they allow basic traffic to flow as expected.
- If ping does not work, confirm that the FortiGate interfaces have Ping enabled and also ensure Ping is enabled in the firewall policies. Otherwise the Ping traffic will be blocked.
- If one protocol does not work but others do work, check the FortiGate firewall policies for that one protocol to ensure it is allowed.
- If there are unexplained connectivity problems, check the local computer to ensure it does not have a software firewall running that may be blocking traffic. MS Windows computers have a firewall running by default that can cause problems.

For additional troubleshooting, see [Troubleshooting VDOMs](#).

Troubleshooting VDOMs

When you're configuring VDOMs you may run into some issues, with your VDOM configuration, your network configuration, or your device setup. This section addresses common problems and specific concerns that an administrator of a VDOM network may have.

This section includes:

- [VDOM admin having problems gaining access](#)
- [FortiGate running very slowly](#)
- [General VDOM tips and troubleshooting](#)

VDOM admin having problems gaining access

With VDOMs configured, administrators have an extra layer of permissions and may have problems accessing their information.

Confirm the admin's VDOM

Each administrator account, other than the super_admin account, is tied to one specific VDOM. That administrator is not able to access any other VDOM. It may be possible they are trying to access the wrong VDOM.

Confirm the VDOM interfaces

An administrator can only access their VDOM through interfaces that are assigned to that VDOM. If interfaces on that VDOM are disabled or unavailable there will be no method of accessing that VDOM by its local administrator. The super_admin will be required to either bring up the interfaces, fix the interfaces, or move another interface to that VDOM to restore access.

Confirm the VDOMs admin access

As with all FortiGate devices, administration access on the VDOM interfaces must be enabled for that VDOM administrators to gain access. For example, if SSH is not enabled, that is not available to administrators.

To enable admin access, the super_admin will go to the **Global > Network > Interfaces** page, and for the interface in question enable the admin access.

FortiGate running very slowly

You may experience a number of problems resulting from your FortiGate being overloaded. These problems may appear as:

- CPU and memory threshold limits exceeded on a continual basis
- Antivirus failopen happening on a regular basis
- dropped traffic or sessions due to lack of resources

These problems are caused by a lack of system resources. There are a number of possible reasons for this.

Too many VDOMs

If you have configured many VDOMs on your system, past the default ten VDOMs, this could easily be your problem.

Each VDOM you create on your FortiGate requires system resources to function - CPU cycles, memory, and disk space. When there are too many VDOMs configured there are not enough resources for operation. This may be a lack of memory in the session table, or no CPU cycles for processing incoming IPS traffic, or even a full disk drive.

Go to **Global > System > VDOM** and see the number of configured VDOMs on your system. If you're running 500 or more VDOMs, you must have a FortiGate 5000 chassis. Otherwise you need to reduce the number of VDOMs on your system to fix the problem. Even if you have the proper hardware, you may encounter noticeably slow throughput if you're using advanced features such as security profiles or deep content inspection with many configured VDOMs.

One or more VDOMs are consuming all the resources

If you have sufficient hardware to support the number of VDOMs you're running, check the global resources on your FortiGate. At a glance it will tell you if you're running out of a particular resource such as sessions, or users. If this is the case, you can then check your VDOMs to see if one particular VDOM is using more than its share of resources. If that is the case you can change the resource settings to allow that VDOM (or those VDOMs) fewer resources and in turn allow the other VDOMs access to those resources.

Too many security features in use

It is likely that reducing the security features in use regardless of number of VDOMs will greatly improve overall system performance and should be considered as an option.

Finally it is possible that your FortiGate configuration is incorrect in some other area, which is using up all your resources. For example, forgetting that you're running a network sniffer on an interface will create significant amounts of traffic that may prevent normal operation.

General VDOM tips and troubleshooting

Besides ping and traceroute, there are additional tools for troubleshooting your VDOM configurations. These include packet sniffing and debugging the packet flow.

Perform a sniffer trace

When troubleshooting networks, it helps to look inside the headers of packets to determine if they are traveling along the route you expect that they are. Packet sniffing can also be called a network tap, packet capture, or logic analyzing.



If your FortiGate has NP interfaces that are offloading traffic, this will change the sniffer trace. Before performing a trace on any NP interfaces, you should disable offloading on those interfaces.

What sniffing packets can tell you

If you're running a constant traffic application such as ping, packet sniffing can tell you if the traffic is reaching the destination, what the port of entry is on the FortiGate, if the ARP resolution is correct, and if the traffic is being sent back to the source as expected.

Sniffing packets can also tell you if the FortiGate is silently dropping packets for reasons such as RPF (Reverse Path Forwarding), also called Anti Spoofing, which prevents an IP packet from being forwarded if its Source IP does not either belong to a locally attached subnet (local interface), or be part of the routing between the FortiGate and another source (static route, RIP, OSPF, BGP). Note that RPF can be disabled by turning on asymmetric routing in the CLI (`config system setting, set asymmetric enable`), however this will disable stateful inspection on the FortiGate and cause many features to be turned off.



If you configure virtual IP addresses on your FortiGate, it will use those addresses in preference to the physical IP addresses. You will notice this when you're sniffing packets because all the traffic will be using the virtual IP addresses. This is due to the ARP update that is sent out when the VIP address is configured.

How to sniff packets

When you're using VDOMs, you must be in a VDOM to access the `diag sniffer` command. At the global level, the command is not available. This is limit the packets only to the ones on your VDOM, and protects the privacy of other VDOM clients.

The general form of the internal FortiOS packet sniffer command is:

```
diag sniffer packet <interface_name> <'filter'> <verbose> <count>
```

To stop the sniffer, type `CTRL+C`.

<interface_name>	The name of the interface to sniff, such as “port1” or “internal”. This can also be “any” to sniff all interfaces.
<'filter'>	What to look for in the information the sniffer reads. <code>none</code> indicates no filtering, and all packets will be displayed as the other arguments indicate. The filter must be inside single quotes (').
<verbose>	The level of verbosity as one of: 1 - print header of packets 2 - print header and data from IP of packets 3 - print header and data from Ethernet of packets
<count>	The number of packets the sniffer reads before stopping. If you don't put a number here, the sniffer will run forever until you stop it with <CTRL C>.

For a simple sniffing example, enter the CLI command `diag sniffer packet port1 none 1 3`. This will display the next 3 packets on the port1 interface using no filtering, and using verbose level 1. At this verbosity level you can see the source IP and port, the destination IP and port, action (such as ack), and sequence numbers.

In the output below, port 443 indicates these are HTTPS packets, and 172.20.120.17 is both sending and receiving traffic.

```
Head_Office_620b # diag sniffer packet port1 none 1 3
interfaces=[port1]
filters=[none]
0.545306 172.20.120.17.52989 -> 172.20.120.141.443: psh 3177924955 ack 1854307757

0.545963 172.20.120.141.443 -> 172.20.120.17.52989: psh 1854307757 ack 3177925808

0.562409 172.20.120.17.52988 -> 172.20.120.141.443: psh 4225311614 ack 3314279933
```

For a more advanced example of packet sniffing, the following commands will report packets on any interface traveling between a computer with the host name of PC1 and the computer with the host name of PC2. With verbosity 4 and above, the sniffer trace will display the interface names where traffic enters or leaves the FortiGate. Remember to stop the sniffer, type CTRL+C. Note that PC1 and PC2 may be VDOMs.

```
FGT# diagnose sniffer packet any "host <PC1> or host <PC2>" 4
```

or

```
FGT# diagnose sniffer packet any "(host <PC1> or host <PC2>) and icmp" 4
```

The following sniffer CLI command includes the ARP protocol in the filter which may be useful to troubleshoot a failure in the ARP resolution (for instance PC2 may be down and not responding to the FortiGate ARP requests).

```
FGT# diagnose sniffer packet any "host <PC1> or host <PC2> or arp" 4
```


Debugging the packet flow

Traffic should come in and leave the VDOM. If you have determined that network traffic is not entering and leaving the VDOM as expected, debug the packet flow.

Debugging can only be performed using CLI commands. Debugging the packet flow requires a number of debug commands to be entered as each one configures part of the debug action, with the final command starting the debug.



If your FortiGate has NP interfaces that are offloading traffic, this will change the packet flow. Before performing the debug on any NP interfaces, you should disable offloading on those interfaces.

The following configuration assumes that PC1 is connected to the internal interface of the FortiGate and has an IP address of 10.11.101.200. PC1 is the host name of the computer.

To debug the packet flow in the CLI, enter the following commands:

```
FGT# diag debug enable
FGT# diag debug flow filter add <PC1>
FGT# diag debug flow show console enable
FGT# diag debug flow trace start 100
FGT# diag debug enable
```

The `start 100` argument in the above list of commands will limit the output to 100 packets from the flow. This is useful for looking at the flow without flooding your log or your display with too much information.

To stop all other debug activities, enter the command:

```
FGT# diag debug flow trace stop
```

The following is an example of debug flow output for traffic that has no matching Firewall Policy, and is in turn blocked by the FortiGate. The denied message indicates the traffic was blocked. Note that even with VDOMs not enabled, vd-root is still shown.

```
id=20085 trace_id=319 func=resolve_ip_tuple_fast line=2825 msg="vd-root received a
packet(proto=6, 192.168.129.136:2854->192.168.96.153:1863) from port3."

id=20085 trace_id=319 func=resolve_ip_tuple line=2924 msg="allocate a new session-
013004ac"

id=20085 trace_id=319 func=vf_ip4_route_input line=1597 msg="find a route: gw-
192.168.150.129 via port1"

id=20085 trace_id=319 func=fw_forward_handler line=248 msg=" Denied by forward policy
check"
```

Virtual FortiOS (Private Cloud Administration Guide)

This document describes how to deploy a FortiGate virtual appliance in several virtualization server environments. This includes how to configure the virtual hardware settings of the virtual appliance.

This document assumes:

- you have already successfully installed the virtualization server on the physical machine,
- you have installed appropriate VM management software on either the physical server or a computer to be used for VM management.

This document does not cover configuration and operation of the virtual appliance after it has been successfully installed and started. For these issues, see the FortiGate Handbook.

This document includes the following sections:

- [Virtual FortiOS overview](#)
- [Deployment example – VMware](#)
- [Deployment example – MS Hyper-V](#)
- [Deployment example – KVM](#)
- [Deployment example – Open Xen](#)
- [Deployment example – Citrix XenServer](#)

What's new in virtual FortiOS 6.0.1

The following list contains new firewall features added in FortiOS 6.0.1. Click on a link to navigate to that section for further information.

- [SDN connector support on page 3035](#)
- [Tagging NSX VMs from FortiGate on page 3028](#)
- [AWS GuardDuty integration on page 3030](#)
- [HA support for Oracle Cloud Infrastructure on page 3036](#)

What's new in virtual FortiOS 6.0

The following list contains new firewall features added in FortiOS 6.0. Click on a link to navigate to that section for further information.

- [FortiOS On-Demand supported platforms on page 3027](#)
- [NSX credentials encryption on page 3028](#)
- [Certificate validation to NSX manager on page 3028](#)
- [FortiGate VMX license status](#)
- [SDN connector addressing on page 3029](#)
- [Support for KVM-based hypervisor in AWS on page 3029](#)

- [HA support for GCP on page 3034](#)
- [HA support for Azure on page 3032](#)

Virtual FortiOS overview

The following topics are included in this section:

FortiGate VM models and licensing

Fortinet offers the FortiGate VM in five virtual appliance models determined by license. When configuring your FortiGate VM, be sure to configure hardware settings within the ranges outlined below. Contact your Fortinet Authorized Reseller for more information.

FortiGate VM model information

Technical Specification	FG-VM00	FG-VM01	FG-VM02	FG-VM04	FG-VM08	FG-VM16	FG-VM32	FG-VMUL
Virtual CPUs (min / max)	1 / 1	1 / 1	1 / 2	1 / 4	1 / 8	1 / 16	1 / 32	1 / unlimited
Virtual Network Interfaces (min / max)	2 / 10							
Virtual Memory (min / max)	1GB / 2GB	1GB / 2GB	1GB / 4GB	1GB / 6GB	1GB / 12GB	1GB / 24GB	1GB / 48GB	1GB / unlimited GB
Virtual Storage (min / max)	32GB / 2TB							
Managed Wireless APs (tunnel mode / global)	32 / 32	32 / 64	256 / 512	256 / 512	1024 / 4096	1024 / 4096	1024 / 4096	1024 / 4096
Virtual Domains (default / max)	1 / 2	10 / 10	10 / 25	10 / 50	10 / 500	10 / 500	10 / 500	10 / 500



There may be times the min/max values can change. An example for this is when the maximum memory for FG-VM00 changed between 5.2 and 5.4 from 1 GB to 1.5 GB. If that is the case, the settings for the VM will have to be manually changed to accommodate the new parameters.

After placing an order for FortiGate VM, a license registration code is sent to the email address used on the order form. Use the registration number provided to register the FortiGate VM with Customer Service & Support and then download the license file. Once the license file is uploaded to the FortiGate VM and validated, your FortiGate VM appliance is fully functional.

FortiGate VM evaluation license

FortiGate VM includes a limited embedded 15-day trial license that supports:

- 1 CPU maximum
- 1024 MB memory maximum
- low encryption only (no HTTPS administrative access)
- all features except FortiGuard updates

You cannot upgrade the firmware, doing so will lock the Web-based Manager until a license is uploaded. Technical support is not included. The trial period begins the first time you start FortiGate VM. After the trial license expires, functionality is disabled until you upload a license file.



The number of Virtual Network Interfaces is not solely dependent on the FortiGate VM. Some virtual environments have their own limitations on the number of interfaces allowed. As an example, if you go to <https://docs.microsoft.com/en-us/azure/virtual-network/virtual-networks-multiple-nics>, you will find that Azure has its own restrictions for VMs, depending on the type of deployment or even the size of the VM.

Registering FortiGate VM

To obtain the FortiGate VM license file you must first register your FortiGate VM with the [Fortinet Support](#) website.

To register your FortiGate VM:

1. Log in to the Customer Service & Support portal using an existing support account or select **Sign Up** to create a new account.
2. In the main page, under **Asset**, select **Register/Renew**.

The **Registration** page opens.

3. Enter the registration code that was emailed to you and select **Register**. A registration form will display.
4. After completing the form, a registration acknowledgment page will appear.
5. Select the **License File Download** link.
6. You will be prompted to save the license file (`.lic`) to your local computer. See "Upload the license file" for instructions on uploading the license file to your FortiGate VM via the Web-based Manager.

Downloading the FortiGate VM deployment package

FortiGate VM deployment packages are included with FortiGate firmware images on the [Customer Service & Support](#) site. First, see the following table to determine the appropriate VM deployment package for your VM platform.

Selecting the correct FortiGate VM deployment package for your VM platform

VM Platform	FortiGate VM Deployment File
Citrix XenServer v5.6sp2, 6.0 and later	FGT_VM64-v500-buildnnnn-FORTINET.out.CitrixXen.zip

VM Platform	FortiGate VM Deployment File
Open Xen v3.4.3, 4.1	FGT_VM64-v500-buildnnnn-FORTINET.out.OpenXen.zip
Microsoft Hyper-V Server 2008R2 and 2012	FGT_VM64-v500-buildnnnn-FORTINET.out.hyperv.zip
KVM (qemu 0.12.1)	FGT_VM64-v500-buildnnnn-FORTINET.out.kvm.zip
VMware ESX 4.0, 4.1 ESXi 4.0/4.1/5.0/5.1/5.5	FGT_VM32-v500-buildnnnn-FORTINET.out.ovf.zip (32-bit) FGT_VM64-v500-buildnnnn-FORTINET.out.ovf.zip

For more information see the FortiGate product datasheet available on the Fortinet web site, https://www.fortinet.com/content/dam/fortinet/assets/data-sheets/FortiGate_VM.pdf.

The firmware images FTP directory is organized by firmware version, major release, and patch release. The firmware images in the directories follow a specific naming convention and each firmware image is specific to the device model. For example, the FGT_VM32-v500-build0151-FORTINET.out.ovf.zip image found in the v5.0 Patch Release 2 directory is specific to the FortiGate VM 32-bit environment.



You can also download the FortiOS Release Notes, FORTINET-FORTIGATE MIB file, FSSO images, and SSL VPN client in this directory. The Fortinet Core MIB file is located in the main FortiGate v5.00 directory.

To download the FortiGate VM deployment package:

1. In the main page of the Customer Service & Support site, select **Download > Firmware Images**.

The **Firmware Images** page opens.

2. In the **Firmware Images** page, select **FortiGate**.
3. Browse to the appropriate directory on the FTP site for the version that you would like to download.
4. Download the appropriate .zip file for your VM server platform.

You can also download the FortiGate Release Notes.

5. Extract the contents of the deployment package to a new file folder.

Deployment package contents

Citrix XenServer

The FORTINET.out.CitrixXen.zip file contains:

- fortios.vhd: the FortiGate VM system hard disk in VHD format
- fortios.xva: binary file containing virtual hardware configuration settings

- in the ovf folder:
 - FortiGate-VM64.ovf: Open Virtualization Format (OVF) template file, containing virtual hardware settings for Xen
 - fortios.vmdk: the FortiGate VM system hard disk in VMDK format
 - datadrive.vmdk: the FortiGate VM log disk in VMDK format

The ovf folder and its contents is an alternative method of installation to the .xva and VHD disk image.

OpenXEN

The FORTINET.out.OpenXen.zip file contains only fortios.qcow2, the FortiGate VM system hard disk in qcow2 format. You will need to manually:

- create a 32GB log disk
- specify the virtual hardware settings

Microsoft Hyper-V

The FORTINET.out.hyperv.zip file contains:

- in the Virtual Hard Disks folder:
 - fortios.vhd: the FortiGate VM system hard disk in VHD format
 - DATADRIIVE.vhd: the FortiGate VM log disk in VHD format
- In the Virtual Machines folder:
 - fortios.xml: XML file containing virtual hardware configuration settings for Hyper-V. This is compatible with Windows Server 2012.
- Snapshots folder: optionally, Hyper-V stores snapshots of the FortiGate VM state here

KVM

The FORTINET.out.kvm.zip contains only fortios.qcow2, the FortiGate VM system hard disk in qcow2 format. You will need to manually:

- create a 32GB log disk
- specify the virtual hardware settings

VMware ESX/ESXi

You will need to create a 32GB log disk.

The FORTINET.out.ovf.zip file contains:

- fortios.vmdk: the FortiGate VM system hard disk in VMDK format
- datadrive.vmdk: the FortiGate VM log disk in VMDK format
- Open Virtualization Format (OVF) template files:
 - FortiGate-VM64.ovf: OVF template based on Intel e1000 NIC driver
 - FortiGate-VM64.hw04.ovf: OVF template file for older (v3.5) VMware ESX server
 - FortiGate-VMxx.hw07_vmxnet2.ovf: OVF template file for VMware vmxnet2 driver
 - FortiGate-VMxx.hw07_vmxnet3.ovf: OVF template file for VMware vmxnet3 driver



Use the VMXNET3 interface (FortiGate-VMxx.hw07_vmxnet3.ovf template) if the virtual appliance will distribute workload to multiple processor cores.

Deploying the FortiGate VM appliance

Prior to deploying the FortiGate VM appliance, the VM platform must be installed and configured so that it is ready to create virtual machines. The installation instructions for FortiGate VM assume that

- You are familiar with the management software and terminology of your VM platform.
- An Internet connection is available for FortiGate VM to contact FortiGuard to validate its license or, for closed environments, a FortiManager can be contacted to validate the FortiGate VM license. See "Validate the FortiGate VM license with FortiManager".

For assistance in deploying FortiGate VM, refer to the deployment chapter in this guide that corresponds to your VMware environment. You might also need to refer to the documentation provided with your VM server. The deployment chapters are presented as examples because for any particular VM server there are multiple ways to create a virtual machine. There are command line tools, APIs, and even alternative graphical user interface tools.

Before you start your FortiGate VM appliance for the first time, you might need to adjust virtual disk sizes and networking settings. The first time you start FortiGate VM, you will have access only through the console window of your VM server environment. After you configure one FortiGate network interface with an IP address and administrative access, you can access the FortiGate VM web-based manager.

After deployment and license validation, you can upgrade your FortiGate VM appliance's firmware by downloading either FGT_VM32-v500-buildnnnn-FORTINET.out (32-bit) or FGT_VM64-v500-buildnnnn-FORTINET.out (64-bit) firmware. Firmware upgrading on a VM is very similar to upgrading firmware on a hardware FortiGate unit.



FortiGate-VM is not part of the FortiGuard Network for the purpose of upgrades.

Performance and optimization

Performance is improved for FortiOS VM platforms by implementing features to improve efficiency and resource utilization.

Interrupt affinity

You can configure interrupt affinity and packet distribution to optimize performance for your VM environment. Interrupt affinity allows you to align interrupts from interfaces to specific CPUs.

Configuring interrupt affinity

Use the following commands to configure interrupt affinity for two 10G interfaces (port2 and port3).

Interrupts from first interface are assigned to core #0 and those from the second interface are assigned to core #1.

```
config system affinity-interrupt
edit 1
```



```
    set interrupt "port2-TxRx-0"
    set affinity-cpumask "0x1"
next
edit 2
    set interrupt "port2-TxRx-1"
    set affinity-cpumask "0x1"
next
edit 3
    set interrupt "port3-TxRx-0"
    set affinity-cpumask "0x2"
next
edit 4
    set interrupt "port3-TxRx-1"
    set affinity-cpumask "0x2"
end
```

Packet distribution

Packet distribution allows you to configure FortiGate-VM to distribute processing to multiple CPUs. Use the following commands to configure packet redistribution to redistribute packets from core #0 and #1 to all other cores.

Configuring packet distribution

The example is based on VM08:

```
config system affinity-packet-redistribution
edit 1
    set interface "port2"
    set affinity-cpumask "0xFC"
next
edit 2
    set interface "port3"
    set affinity-cpumask "0xFC"
end
```

Other virtual FortiOS products

Just like a VM installed on a platform located on a physical computer, these are instances of the FortiOS firmware installed in virtual environments.

SDN environments

FortiOS-VM is now supported in a Number of SDN (Software Defined Network) environments through the use of a SDN connector. The software can be found on the [Fortinet Service and Support site](#). Documentation can be found on the [Fortinet Documentation site](#) under the product heading of Fortinet Connectors.

The products include:

FortiGate-VMX

Admin guide can be found at <https://docs.fortinet.com/d/fortigate-vmx-install-v.2>.

FortiOS On-Demand

Admin guide can be found at <https://docs.fortinet.com/d/fortigate-fortinetvm-on-demand-administration-guide>.

FortiOS On-Demand supported platforms

Connectors:

FortiADC Connector for Cisco ACI

- Guides found at <https://docs.fortinet.com/fortinet-connectors/admin-guides>.
- Versions available:
 - 1.0.0
 - 1.2.0
 - 1.3.0

FortiADC Connector for Cisco ACI

- Guides found at <https://docs.fortinet.com/fortinet-connectors/admin-guides>.
- Versions available:
 - 1.0.0
 - 1.2.0
 - 1.3.0

FortiGate Connector for HPE VAN SDN

- Guides found at <https://docs.fortinet.com/fortinet-connectors/admin-guides>.
- Versions available:
 - 1.0.5
 - 1.1.0

FortiGate Connector for OpenStack ML2

- Guides found at <https://docs.fortinet.com/fortinet-connectors/admin-guides>.
- Versions available:
 - 1.1

FortiManager Connector for Nuage Networks

- Guides found at <https://docs.fortinet.com/fortinet-connectors/admin-guides>.
- Versions available:
 - 1.0

Fortinet SDN Connector for Nuage VSP

Instruction can be found at <http://cookbook.fortinet.com/nuage-vsp/>

Additional information relating to connectors

While the product documentation concerns itself with the connector software, there is also settings within FortiOS that relates to the usage of those connectors.

NSX connector support

There is support for NSX connector to upgrade to SDN connectors.

Change to config system sdn-connector:

```
config nsx setting
```

Change to config firewall address:

```
set type nsx

set type dynamic
set sdn nsx
```



When using the sdn nsx setting, the user should also use the nsx rest-api password.

Tagging NSX VMs from FortiGate

There is a requirement to have VMs tagged for the use of features such as IP detection. The tagging of VMs in an NSX environment is performed from the FortiGate by using an `execute` command.

The command performs the following functions when run:

- An API call to NSX to find the `vm_id` based on the IP address.
- An API call to resolve the security tag name to `tag_id`
- An API call to add the VM to the security tag

Command syntax:

```
execute sdn tag nsx <ip> <tagname> <enable or disable spoofguard>
```

Option	Description
<code>ip</code>	IP address of the VM
<code>tagname</code>	Name of tag for the VM(s)
<code>enable or disable spoofguard</code>	<ul style="list-style-type: none">• 1 is used to enable SpoofGuard• 0 is used to disable SpoofGuard

NSX credentials encryption

NSX requires AES256 as a minimum level of encryption. The credentials for NSX are stored using AES256 encryption in preparation for being transmitted over the network for NSX authentication.

Certificate validation to NSX manager

The first time FortiGate-SVM connects to the NSX Manager (when a **Service Manager** is added), you can save the service manager's certificate as a fingerprint in the FortiGate configuration. After this is done, the FortiGate-SVM can validate the connection by matching the certificate of the NSX Manager against the saved fingerprint in the configuration every time that:

- the FortiGate-SVM connects to the NSX Manager
- the NSX Manager connects to the FortiGate-SVM

CLI settings:

This setting is available only in the CLI.

```
config system sdn-connector
edit <connector-name>
set type nsx
set nsx-cert-fingerprint <string value>
end
```



`nsx-cert-fingerprint` is a hidden entry. If you are in the `nsx` context and type the command `set ?` to get a listing of the available options, it will not show up.

Removal of fingerprint

If the NSX Manager service is deleted from the FortiGate-SVM, to avoid potential security issues, the fingerprint is automatically deleted.

SDN connector addressing

When setting the address for an SDN connection, the value can be a fully qualified domain name (FQDN) as well as an IP address. The setting that is currently called `server` was once called `server-ip`.

CLI

```
config system sdn-connector
edit <example>
set server [<ip address>|<fqdn address>]
end
```

Public cloud environments

Unlike SDNs where the user has administrative control over the virtual environment that the FortiOS is being placed into, the Public Cloud services are run and managed by 3rd party companies with their own methods and rules of provisioning that need to be followed to install the firmware into the environment.

Amazon Web Services (AWS)

Online documentation can be found at <http://cookbook.fortinet.com/amazon-web-services-aws/>

Additional information for AWS

While the online documentation is primarily for the installation of the FortiOS into the virtual environment, periodically improvements are made to the firmware that affects its interaction with the environment but isn't really covered by installation instructions or information has changed in regards to those instructions but they may not have all been updated yet.

Support for KVM-based hypervisor in AWS

FortiGate-VM can be provisioned in AWS C5 instances on the marketplace (BYOL/OnDemand). These instances are based on a home-brewed version of KVM that AWS started using. Systems deployed in C5 will see better

performance compared to C3/C4 instances. Deploying FortiGate-VM in these instances requires no configuration on the part of the user.

FortiGate-VM firmware can be placed in these instances because of:

- An NVME driver that is required to run C5 instances. It gets the correct partition name dependencies if an NVME device is being used.
- Removing the `xenstore` checking requirement in the AWS setup daemon that is not required in non-Xen based instances.

HA

AWS supports the use of HA.

This includes two parts:

1. HA with unicast heartbeat traffic.
2. AWS API supports to move secondary IPs and update routing tables.

CLI:

Unicast HA config

```
config system ha
    unicast-hb {enable|disable}
    unicast-hb-peerip <Unicast Heartbeat Peer IP>
end
```

SDN connector - AWS

Improvements have been made in the support of FortiGate-VM integrating into the AWS environment.

1. `config aws setting` has been moved to the context of `config system sdn-connector`.

```
config system sdn-connector
    edit <string>
        set access-key <AWS access key ID>
        set secret-key <AWS secret access key>
        set region <AWS region name>
        set vpc-id <AWS VPC ID>
    end
```
2. Update to the GUI SDN connector edit page that supports allowing configuration of the following fields:
 - AWS access key ID
 - AWS secret access key
 - AWS region name
 - AWS VPC ID
 - Update Interval
3. Change to address edit page to allow configuration of the **Filter** field for Dynamic AWS address.
4. Update to the dynamic address monitor API to get resolved address list for dynamic AWS addresses.

AWS GuardDuty integration

AWS GuardDuty is a managed threat detection service that monitors malicious or unauthorized behaviors/activities related to AWS resources. GuardDuty provides visibility of logs called findings, and Fortinet

provides a Lambda script that populates a list of malicious IP addresses then stores it in an S3 location. FortiGate can then be configured to point to the location as the external feed of threat vectors.

To use this feature, you must subscribe to GuardDuty, CloudWatch, and S3.



Fortinet-provided Lambda scripts are not supported within regular Fortinet technical support scope. For questions related to the scripts, contact awssales@fortinet.com.

GuardDuty findings give visibility on the following:

- **Severity:** High/medium/low (associated with scores)
- **Where it occurred:** Region, resource ID, account ID
- **When:** Last seen date/time
- Count
- **Detailed information can include:**
 - **Affected resource:** type/instance ID/image ID/port/resource type/image description/launch time/tags/network interfaces (public IP, private IP, subnet ID, VPCID, security groups)
 - **Action:** type/connection direction
 - Actor
 - Additional information

To configure the integration:

- Subscribe and enable GuardDuty on AWS. When findings occur, they are pushed to CloudWatch.
- CloudWatch events trigger the Lambda script for automated actions.
- If one of the following criteria is met:
 - Connected direction is inbound, the finding contains an IP address, and the severity is greater than the minimum score (configurable)
 - Connected direction is unknown, the finding contains an IP address and matches certain known threat lists (such as ProofPoint) that GuardDuty identifies, and the severity is greater than the minimum score

The IP address is considered black and is appended to a file located in the S3 bucket/directory.

- FortiGate queries the file as the external source of blacklisted IP addresses. The following is an example configuration:

Edit External Resource

Type: FortiGuard Category | **Firewall IP Address** | Domain Name

Name: GuardDuty

URI of external resource: https://s3.us-east-2.amazonaws.com/ip

Refresh Rate: 5

Comments: 0/255

Last Update: Not updated

Status: ☒

OK Cancel

The configuration can be done in the CLI as follows:

```
config system external-resource
  edit "GuardDuty"
    set type address
    set resource "https://s3.us-east-2.amazonaws.com/ip-blacklist/ip.txt"
  next
end
```

You can then use the gathered IP addresses as criteria to protect the network.

Azure

HA support for Azure

FortiOS supports the use of active/passive HA, similar to that for Amazon Web Services (AWS) in an Azure environment.

New CLI options are available to configure with Azure:

- The `system sdn-connector` contains a connector type for Azure environments called `azure`.
- The `nic` option enables you to configure the Azure network interface.
- The `route-table` option enables you to configure the Azure route table.

Example:

```
config system sdn-connector
  edit "azd"
    set type azure
  config nic
    edit nic-eth2
      config ip
        edit ipconfig1
          set public-ip p3
        next
      end
    next
  edit nic-eth3
    config ip
```

```

        edit ipconfig1
            set public-ip p2
        next
    end
next
end
config route-table
    edit 1
        config route
            edit 12
                set next-hop 111.233.222.233
            next
        end
    next
end

```

In addition to the options for HA already specified, there are settings specific to Azure connectivity within the `system sdn-connector` context once the `type` has been set to `azure`.

Option	Description
tenant-id	Azure tenant ID (directory ID).
subscription-id	Azure subscription ID.
client-id	Azure client ID (application ID).
client-secret	Azure client secret (application key).
resource-group	Azure resource group.
azure-region	Global/China Azure Region.
update-interval	Dynamic object update interval in seconds (0 - 3600, 0 means disabled, default = 60).

Online documentation can be found at <http://cookbook.fortinet.com/microsoft-azure/>

Google Cloud Platform (GCP)

The following FortiGate-VM models will be supported on Google Cloud Platform:

- FG-VM01
- FG-VM02
- FG-VM04
- FG-VM08



FG-WM00 is not supported.

Since GCP use netmask 32, static route must be configured on GCP VPC, instead of FGT.

Licenses will be interchangeable between platforms. A FG-VM04 license that functions in a VMware or Citrix environment can be also used in the GCP environment as well.

While an .out file will be necessary for upgrading, full downloadable images will not be needed for initial installation of the solution. GCP consists of preexisting images that can be checked out of their library and deployed instantly. A difference between this environment and enterprise virtualization platforms is that machine size can never change. An n1-standard-4 has exactly 15 GB of RAM and 4 vCPUs. This can never be changed or edited by the end user or administrator.

The currently available GCP instances we are looking to support are as follows (these will/could change as vNIC values reveal themselves):

FG-VM	Equates to Instance Type	vCPU	RAM	Disks
FG-VM01-GC	n1-standard-1	1	3.75GB	16 (32 in Beta)
FG-VM02-GC	n1-standard-2	2	7.50 GB	16 (64 in Beta)
FG-VM04-GC	n1-standard-4	4	15 GB	16 (64 in Beta)
FG-VM08-GC	n1-standard-8	8	30 GB	16 (128 in Beta)
FG-VM16-GC	n1-standard-16	16	60 GB	16 (128 in Beta)
FG-VM32-GC	n1-standard-32	32	88 GB	16 (128 in Beta)
FG-VMUL-GC	any of the above and any new that could be created.			

Additional information from GCP: <https://cloud.google.com/compute/docs/images/building-custom-os>.

HA support for GCP

FortiOS supports the use of active/passive HA, using a unicast heartbeat, in Google Cloud Platform (GCP) similar to the HA support for Amazon Web Services (AWS).

This support includes:

- HA in the same region within a project
- HA session and configuration synchronization
- Automatic failover to the passive unit (time dependent on configuration), using the built-in API call to GCP for changing the routing path from the active unit to the passive unit.

There is a new connector type, `gcp`, that has been added to the `type` option for `system sdn-connector` as well as a setting for the mask of the unicast heartbeat in `system ha`.

GCP type example:

```
config system sdn-connector
  edit gcp
    set type gcp
  config external-ip
    edit gundam-public
  next
```

```

        end
    config route
        edit gundam-route
        next
    end
next
end

```

Unicast example:

```

config system ha
    set group-id 20
    set group-name "cluster1"
    set mode a-p
    set hbdev "port3" 50
    set ha-mgmt-status enable
    config ha-mgmt-interfaces
        edit 1
            set interface "port4"
            set gateway 172.16.100.1
        next
    end
    set override disable
    set priority 20
    set unicast-hb enable
    set unicast-hb-peerip 172.16.201.3
    set unicast-hb-netmask 255.255.255.0
end

```

Oracle Cloud Infrastructure (OCI)

SDN connector support

Support has been added to enable SDN connector support for FortiGate in Oracle Cloud Infrastructure. This connector makes calls to the OCI API to get information about VMs running in the OCI cloud.

The connector is configured by using the commands:

```

config sys sdn
    edit <name of sdn controller>
        set status enable
        set type oci
        set tenant-id <id>
        set user-id <user id>
        set compartment-id <compartment id>
        set oci-region <OCI server region>
        set oci-cert <OCI certificate>
        set update-interval <interval time>
        set filter <filter argument>
    end

```

The following OCI filters are supported:

- vm_name=<vm name>
- instance_id=<instance id>
- tag.<key>=<value>
- definedtag.<namespace>.<key>=<value>



The values for the filters can have spaces in them as long as the value is inclosed with quotation marks.

HA support for Oracle Cloud Infrastructure

Oracle Cloud Infrastructure (OCI) supports HA for FortiGate-VM. This support does not apply to OCI-Classical or Oracle Public Cloud (OPC).

OCI supports the use of a secondary IP on the virtual network interface (VNIC). This secondary IP is used as the IP of the FortiGate interface instead of the primary IP address, as well as the next hop in the route. When an HA failover occurs, the secondary IP address on the VNIC is changed from master to slave so that the forwarded traffic can go through the new master.

Deployment example – VMware

Once you have downloaded the FGT_VMxx-v5-build0xxx-FORTINET.out.ovf file from <http://support.fortinet.com> and extracted the package contents to a folder on your local computer, you can use the vSphere client to create the virtual machine from the deployment package OVF template.

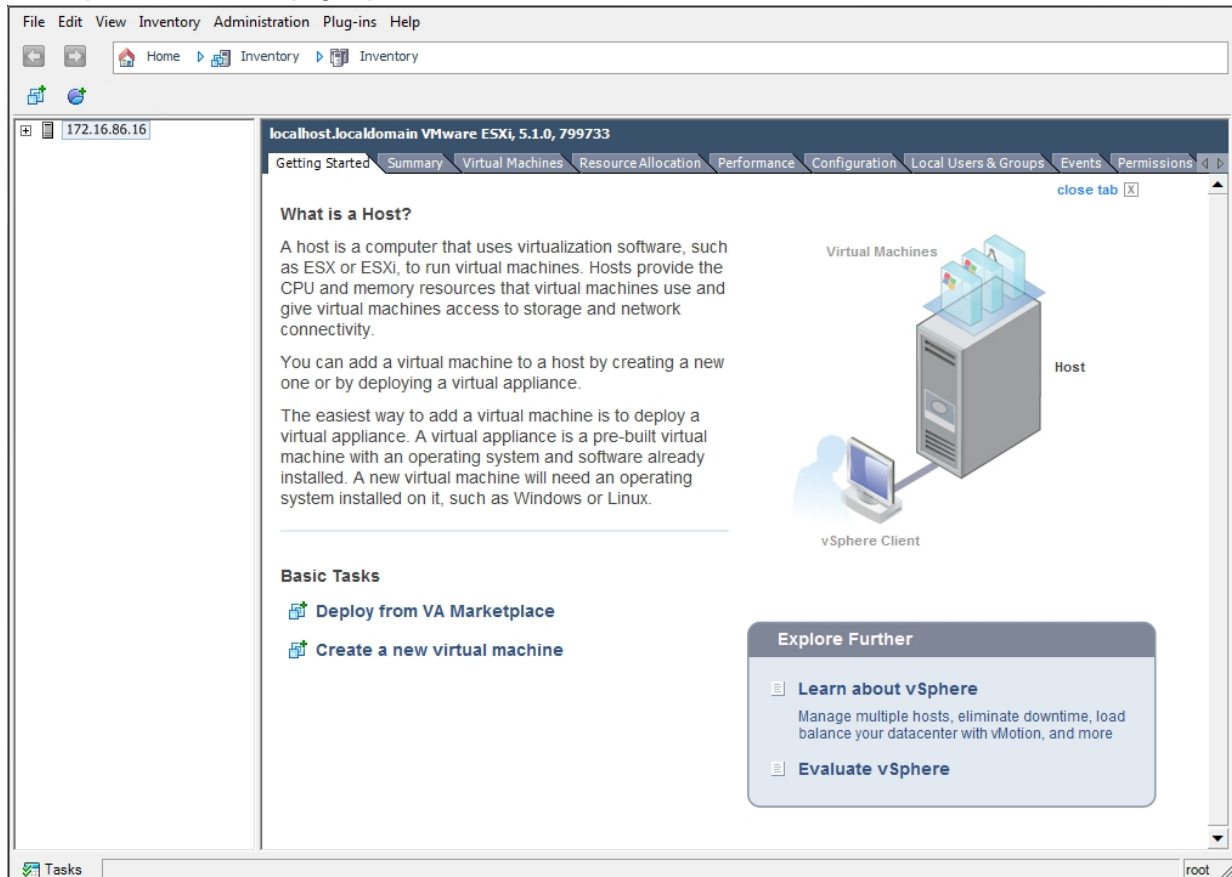
The following topics are included in this section:

Open the FortiGate VM OVF file with the vSphere client

To deploy the FortiGate VM OVF template:

1. Launch the VMware vSphere client, enter the IP address or host name of your server, enter your user name and password and select **Login**.

The vSphere client home page opens.



2. Select **File > Deploy OVF Template** to launch the OVF Template wizard.

The OVF Template **Source** page opens.

3. Select the source location of the OVF file. Select **Browse** and locate the OVF file on your computer. Select **Next** to continue.

The OVF Template **Details** page opens.

OVF Template Details
Verify OVF template details.

[Source](#)

OVF Template Details
End User License Agreement
Name and Location
Disk Format
Network Mapping
Ready to Complete

Product:	Fortigate-VM
Version:	
Vendor:	
Publisher:	No certificate present
Download size:	30.4 MB
Size on disk:	Unknown (thin provisioned) 32.0 GB (thick provisioned)
Description:	FortiGate Virtual Appliance by Fortinet Technologies Inc. (http://www.fortinet.com)

Help < Back Next > Cancel

4. Verify the OVF template details. This page details the product name, download size, size on disk, and description. Select **Next** to continue.

The OVF Template **End User License Agreement** page opens.

End User License Agreement
Accept the end user license agreements.

[Source](#)
[OVF Template Details](#)
End User License Agreement
Name and Location
Disk Format
Network Mapping
Ready to Complete

End User License Agreement for FortiGate Virtual Appliance

NOTICE TO ALL USERS: PLEASE READ THE TERMS AND CONDITIONS OF THE LICENSE AGREEMENT CAREFULLY. FORTINET, INC. IS WILLING TO LICENSE THIS SOFTWARE TO YOU ONLY ON THE CONDITION THAT YOU ACCEPT ALL OF THE TERMS OF THIS LICENSE AGREEMENT. BY CLICKING THE ACCEPT BUTTON OR INSTALLING THE SOFTWARE, YOU (EITHER AN INDIVIDUAL OR A SINGLE ENTITY) AGREE THAT THIS AGREEMENT IS ENFORCEABLE LIKE ANY WRITTEN CONTRACT SIGNED BY YOU. IF YOU DO NOT AGREE, CLICK ON THE BUTTON THAT INDICATES THAT YOU DO NOT ACCEPT THE TERMS OF THIS LICENSE AGREEMENT AND DO NOT INSTALL THE SOFTWARE. IF YOU PURCHASED THE SOFTWARE ON TANGIBLE MEDIA (e.g., CD) WITHOUT THE OPPORTUNITY TO REVIEW THIS LICENSE AND YOU DO NOT ACCEPT THIS LICENSE AGREEMENT, YOU MAY OBTAIN A REFUND OF THE AMOUNT YOU ORIGINALLY PAID IF YOU: (A) DO NOT USE THE SOFTWARE AND (B) RETURN IT, WITH PROOF OF PAYMENT, WITHIN THIRTY (30) DAYS OF THE PURCHASE DATE TO THE LOCATION FROM WHICH IT WAS OBTAINED.

This End User License Agreement (EULA) is an agreement between you and Fortinet, Inc. ("Fortinet"), which governs your use of this software product. A software license and a license key or "unlock code" ("Software License"), issued to a designated user only by Fortinet or its authorized agents, is required for each computer on which the Software is loaded.

Definitions: (a) "Software" means (a) all means (a) all of the contents of the files, disk(s), CD-ROM(s) or other media (including electronic media) with which this Agreement is provided or such contents as are hosted by Fortinet or its distributors, resellers, OEM/MSP partners, or other business partners (collectively "Authorized Partner(s)"), including but not limited to (i) Fortinet or third party computer information or software; (ii) related explanatory materials in printed, electronic, or online form ("Documentation"); and (b) upgrades, modified or subsequent versions and updates (collectively "Updates"), and Software, if any, licensed to you by Fortinet or an

Accept

Help

< Back Next > Cancel

5. Read the end user license agreement for FortiGate VM. Select **Accept** and then select **Next** to continue.

The OVF Template **Name and Location** page opens.

Name and Location
Specify a name and location for the deployed template

[Source](#)
[OVF Template Details](#)
[End User License Agreement](#)
Name and Location
[Disk Format](#)
[Network Mapping](#)
[Ready to Complete](#)

Name:
Fortigate-VM-01
The name can contain up to 80 characters and it must be unique within the inventory folder.

Help < Back Next > Cancel

6. Enter a name for this OVF template. The name can contain up to 80 characters and it must be unique within the inventory folder. Select **Next** to continue.

The OVF Template **Disk Format** page opens.

Disk Format
In which format do you want to store the virtual disks?

Source
OVF Template Details
End User License Agreement
Name and Location
Disk Format
Network Mapping
Ready to Complete

Datastore: datastore 1
Available space (GB): 394.6

☒ Thick Provision Lazy Zeroed
☐ Thick Provision Eager Zeroed
☐ Thin Provision

Help < Back Next > Cancel

7. Select one of the following:
 - **Thick Provision Lazy Zeroed:** Allocates the disk space statically (no other volumes can take the space), but does not write zeros to the blocks until the first write takes place to that block during runtime (which includes a full disk format).
 - **Thick Provision Eager Zeroed:** Allocates the disk space statically (no other volumes can take the space), and writes zeros to all the blocks.
 - **Thin Provision:** Allocates the disk space only when a write occurs to a block, but the total volume size is reported by VMFS to the OS. Other volumes can take the remaining space. This allows you to float space between your servers, and expand your storage when your size monitoring indicates there is a problem. Note that once a Thin Provisioned block is allocated, it remains on the volume regardless if you have deleted data, etc.
8. Select **Next** to continue.

The OVF Template **Network Mapping** page opens.

Network Mapping
What networks should the deployed template use?

[Source](#)
[OVF Template Details](#)
[End User License Agreement](#)
[Name and Location](#)
[Disk Format](#)
Network Mapping
Ready to Complete

Map the networks used in this OVF template to networks in your inventory

Source Networks	Destination Networks
Network 1	VM Network
Network 2	VM Network
Network 3	VM Network
Network 4	VM Network
Network 5	VM Network
Network 6	VM Network
Network 7	VM Network

Description:
The VM Network network

Warning: Multiple source networks are mapped to the host network: VM Network

[Help](#) [< Back](#) [Next >](#) [Cancel](#)

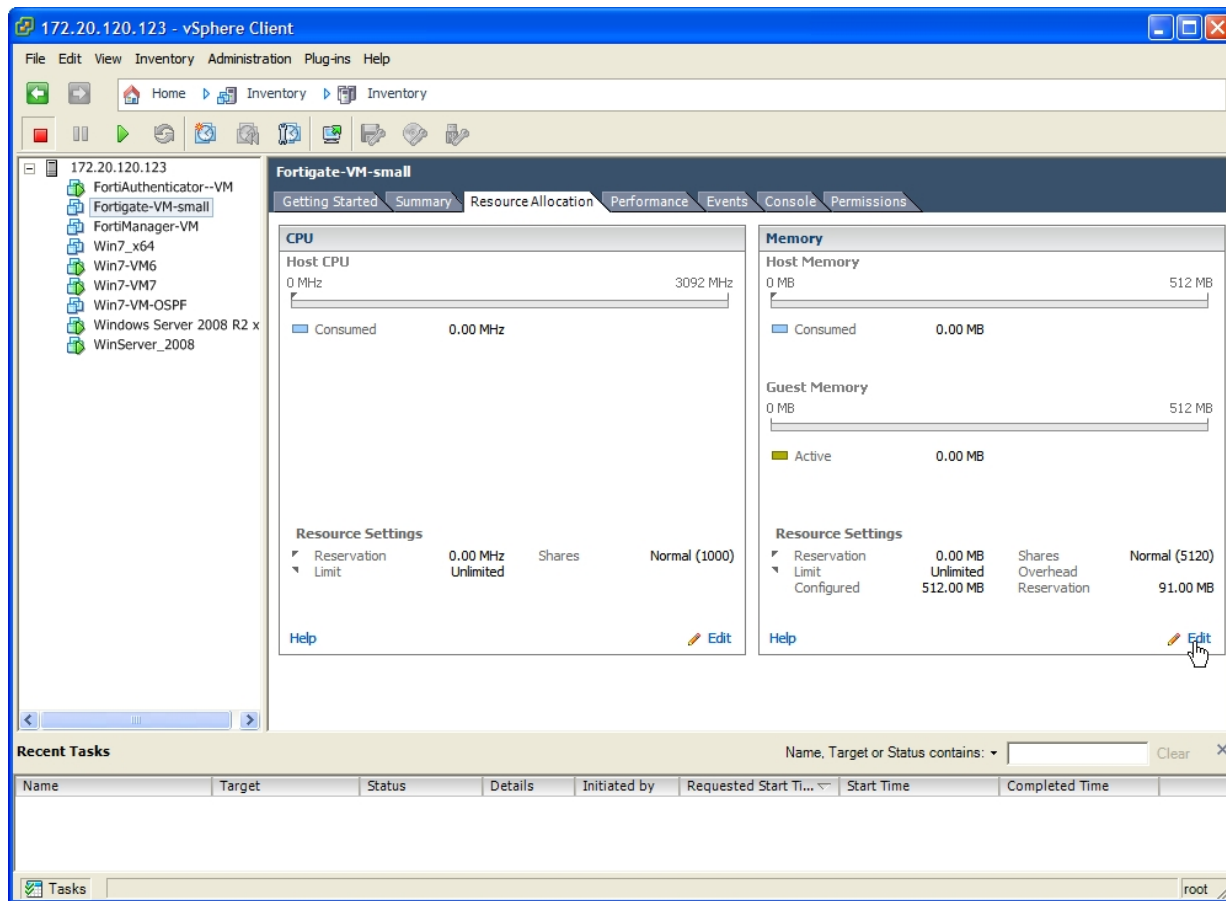
- Map the networks used in this OVF template to networks in your inventory. Network 1 maps to port1 of the FortiGate VM. You must set the destination network for this entry to access the device console. Select **Next** to continue.

The OVF Template **Ready to Complete** page opens.

- Review the template configuration. Make sure that **Power on after deployment** is not enabled. You might need to configure the FortiGate VM hardware settings prior to powering on the FortiGate VM.
- Select **Finish** to deploy the OVF template. You will receive a **Deployment Completed Successfully** dialog box once the FortiGate VM OVF template wizard has finished.

Configure FortiGate VM hardware settings

Before powering on your FortiGate VM you must configure the virtual memory, virtual CPU, and virtual disk configuration to match your FortiGate VM license.



Transparent mode VMware configuration

If you want to use your FortiGate-VM in transparent mode, your VMware server's virtual switches must operate in promiscuous mode. This permits these interfaces to receive traffic that will pass through the FortiGate unit but was not addressed to the FortiGate unit.

In VMware, promiscuous mode must be explicitly enabled:

1. In the vSphere client, select your VMware server in the left pane and then select the **Configuration** tab in the right pane.
2. In **Hardware**, select **Networking**.
3. Select **Properties** of vSwitch0.
4. In the **Properties** window left pane, select **vSwitch** and then select **Edit**.
5. Select the **Security** tab, set **Promiscuous Mode** to **Accept**, then select **OK**.
6. Select **Close**.
7. Repeat steps 3 through 6 for other vSwitches that your transparent mode FortiGate-VM uses.

High availability VMware configuration

If you want to combine two or more FortiGate-VM instances into a FortiGate Clustering Protocol (FGCP) High Availability (HA) cluster the VMware server's virtual switches used to connect the heartbeat interfaces must operate in promiscuous mode. This permits HA heartbeat communication between the heartbeat interfaces. HA

heartbeat packets are non-TCP packets that use Ethertype values 0x8890, 0x8891, and 0x8890. The FGCP uses link-local IPv4 addresses in the 169.254.0.x range for HA heartbeat interface IP addresses.

To enable promiscuous mode in VMware:

1. In the vSphere client, select your VMware server in the left pane and then select the **Configuration** tab in the right pane.
2. In **Hardware**, select **Networking**.
3. Select **Properties** of a virtual switch used to connect heartbeat interfaces.
4. In the **Properties** window left pane, select **vSwitch** and then select **Edit**.
5. Select the **Security** tab, set **Promiscuous Mode** to **Accept**, then select **OK**.
6. Select **Close**.

You must also set the virtual switches connected to other FortiGate interfaces to allow MAC address changes and to accept forged transmits. This is required because the FGCP sets virtual MAC addresses for all FortiGate interfaces and the same interfaces on the different VM instances in the cluster will have the same virtual MAC addresses.

To make the required changes in VMware:

1. In the vSphere client, select your VMware server in the left pane and then select the **Configuration** tab in the right pane.
2. In **Hardware**, select **Networking**.
3. Select **Properties** of a virtual switch used to connect FortiGate VM interfaces.
4. Set **MAC Address Changes** to **Accept**.
5. Set **Forged Transmits** to **Accept**.

Power on your FortiGate VM

You can now proceed to power on your FortiGate VM. There are several ways to do this:

- Select the name of the FortiGate VM you deployed in the inventory list and select **Power on the virtual machine** in the **Getting Started** tab.
- In the inventory list, right-click the name of the FortiGate VM you deployed, and select **Power > Power On**.
- Select the name of the FortiGate VM you deployed in the inventory list. Click the **Power On** button on the toolbar.

Select the Console tab to view the console. To enter text, you must click in the console pane. The mouse is then captured and cannot leave the console screen. As the FortiGate console is text-only, no mouse pointer is visible. To release the mouse, press Ctrl-Alt.

Deployment example – MS Hyper-V

Once you have downloaded the FGT_VMxx_HV-v5-build0xxx-FORTINET.out.hyperv.zip file and extracted the package contents to a folder on your Microsoft server, you can deploy the VHD package to your Microsoft Hyper-V environment.

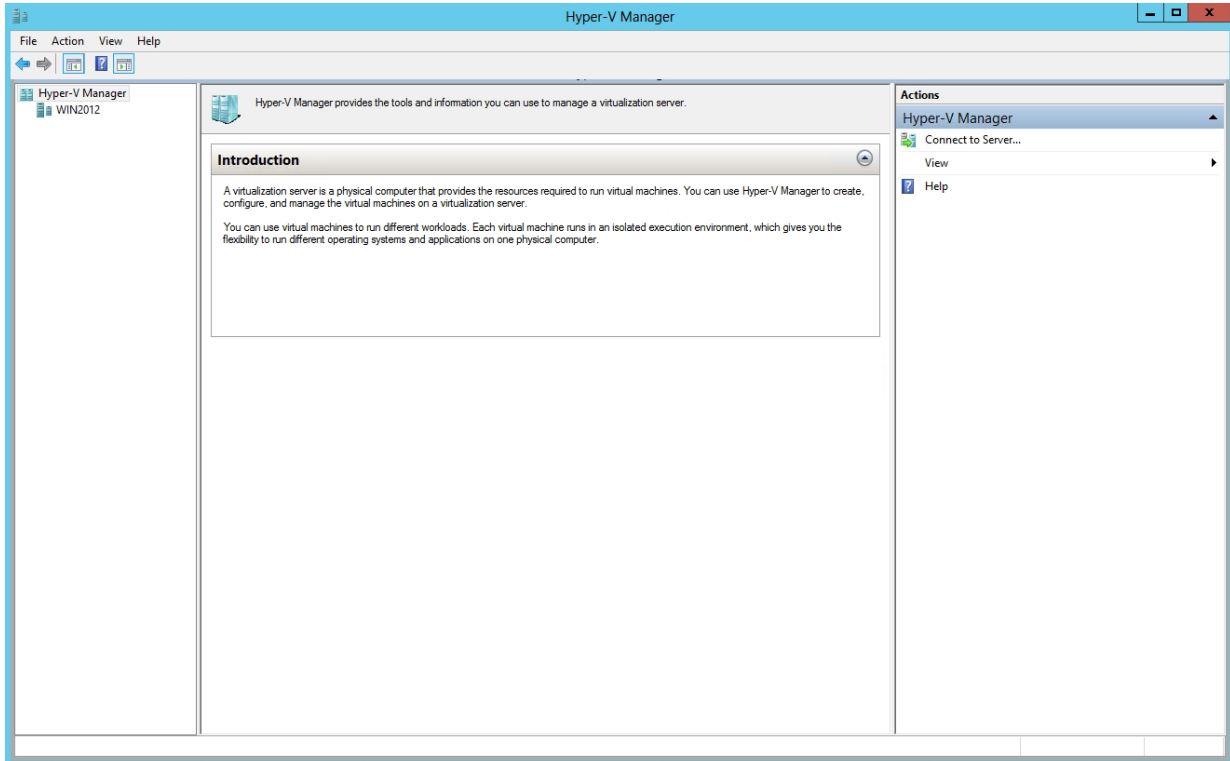
The following topics are included in this section:

Create the FortiGate VM virtual machine

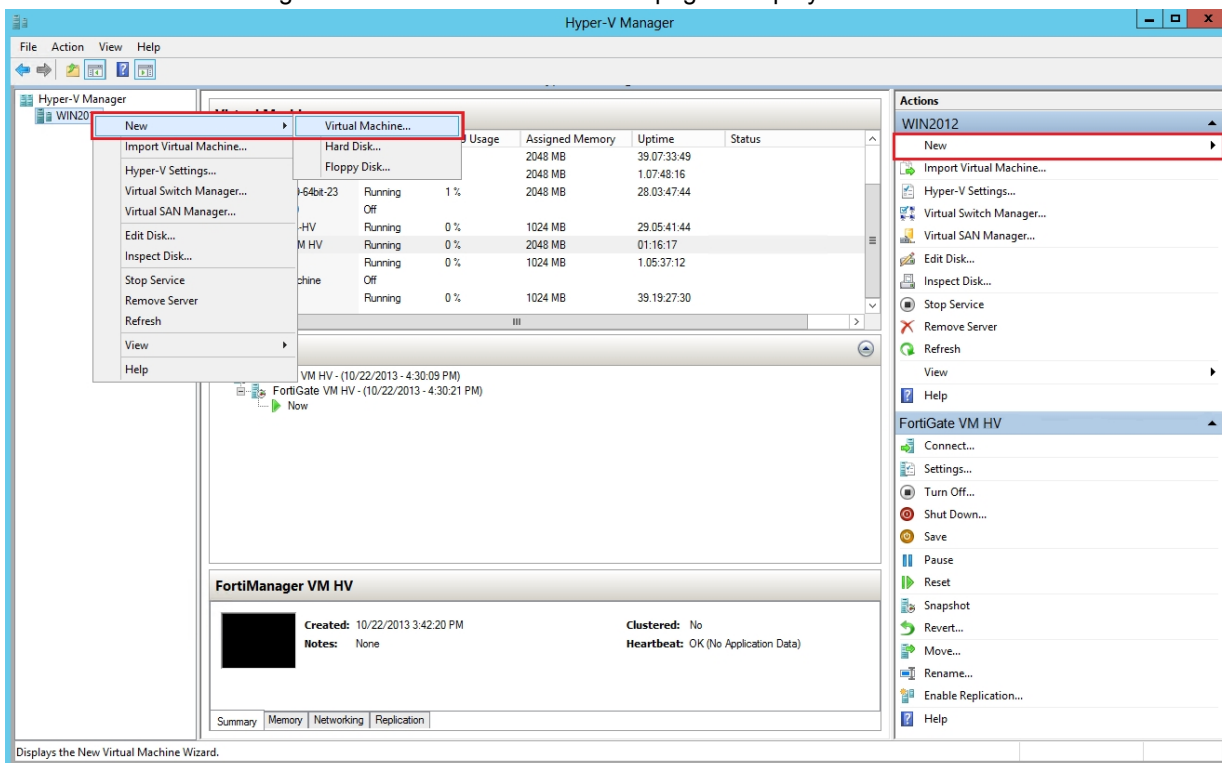
To create the FortiGate VM virtual machine:

1. Launch the Hyper-V Manager in your Microsoft server.

The **Hyper-V Manager** home page opens.

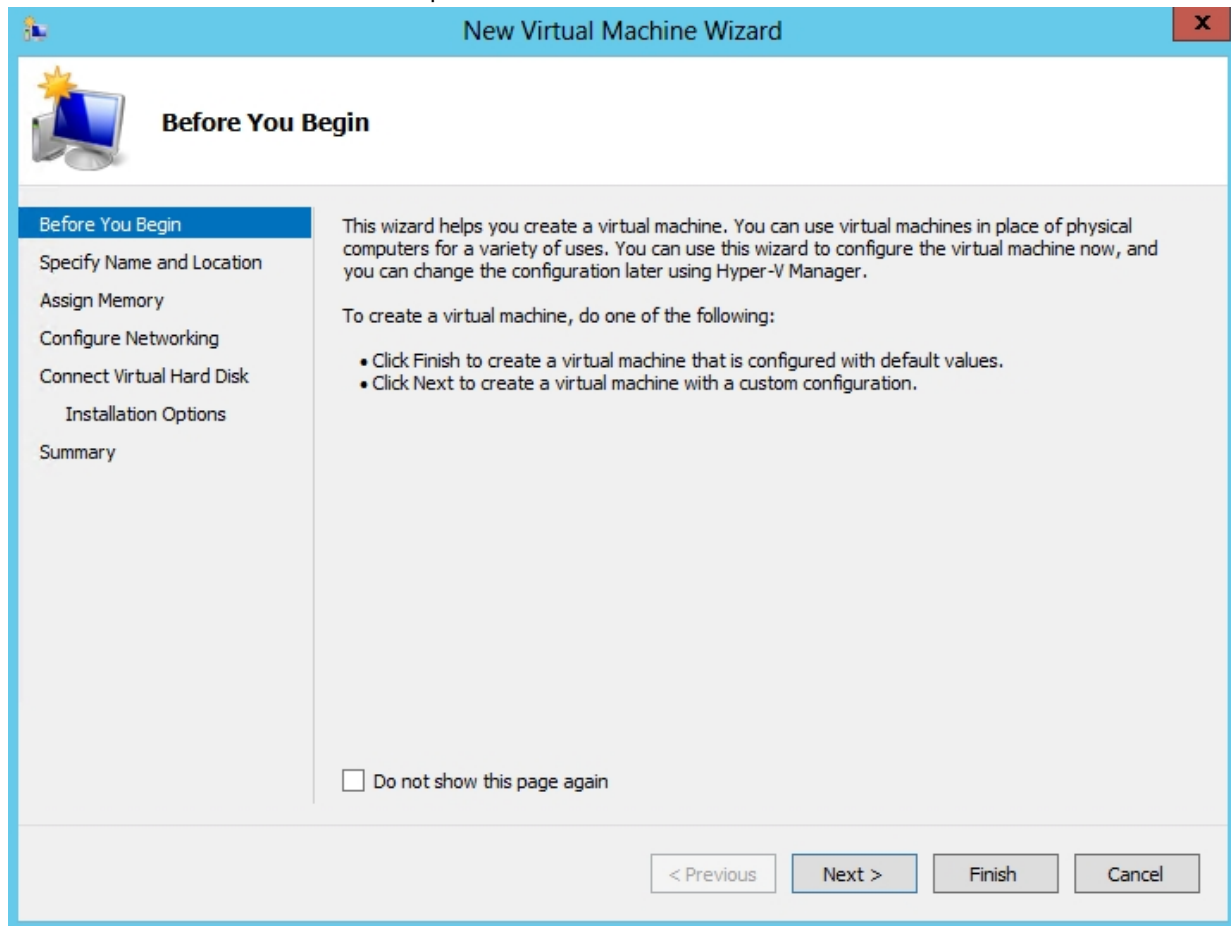


2. Select the server in the right-tree menu. The server details page is displayed.



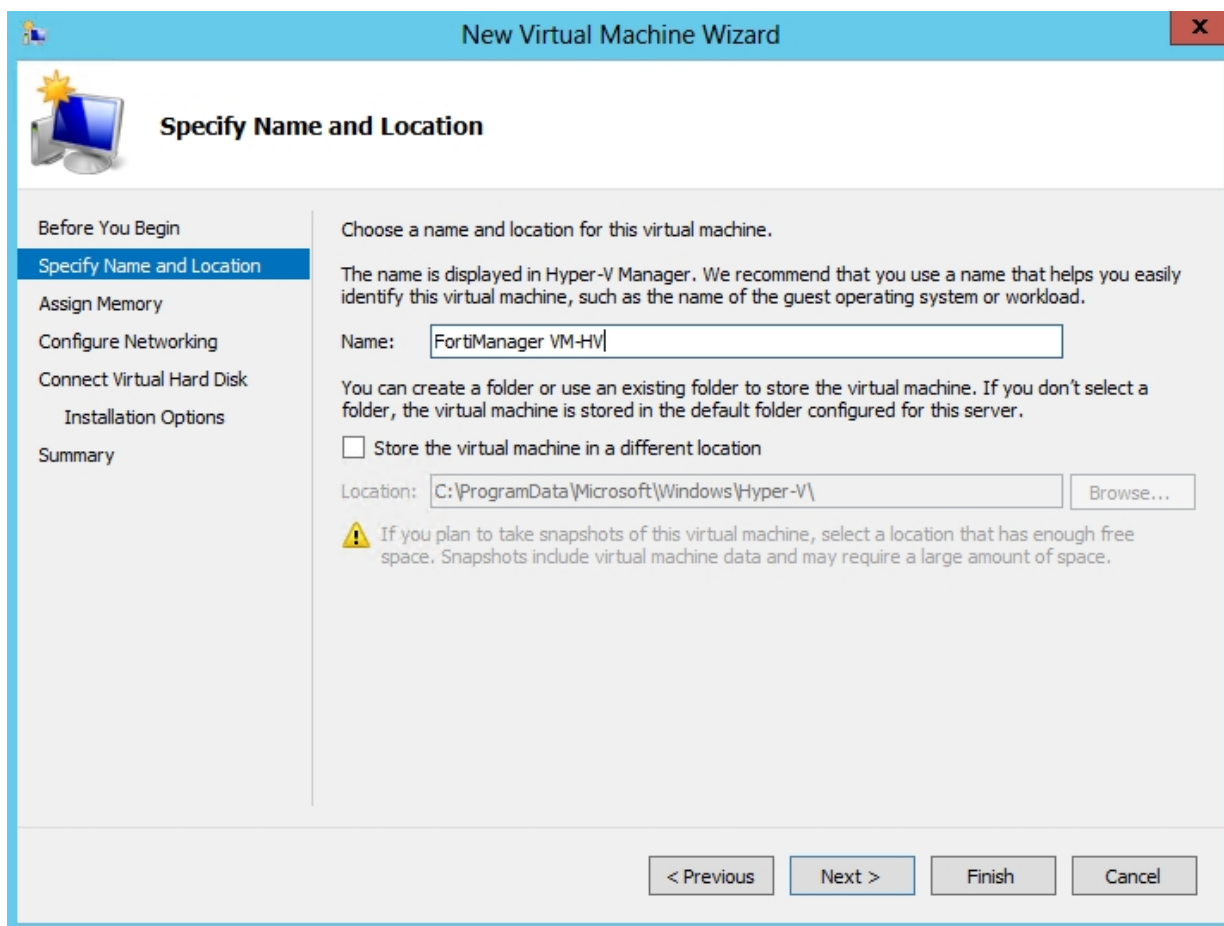
3. Right-click the server and select **New** and select **Virtual Machine** from the menu. Optionally, in the **Actions** menu, select **New** and select **Virtual Machine** from the menu.

The **New Virtual Machine Wizard** opens.



4. Select **Next** to create a virtual machine with a custom configuration.

The **Specify Name and Location** page is displayed.



The screenshot shows the 'New Virtual Machine Wizard' window with the 'Specify Name and Location' step selected in the left-hand navigation pane. The main area contains instructions for naming and locating the virtual machine. The 'Name' field is populated with 'FortiManager VM-HV'. The 'Location' field shows the default path 'C:\ProgramData\Microsoft\Windows\Hyper-V\'. A warning icon and text advise selecting a location with sufficient free space for snapshots. Navigation buttons at the bottom include '< Previous', 'Next >', 'Finish', and 'Cancel'.

New Virtual Machine Wizard

Specify Name and Location

Before You Begin
Specify Name and Location
Assign Memory
Configure Networking
Connect Virtual Hard Disk
Installation Options
Summary

Choose a name and location for this virtual machine.


The name is displayed in Hyper-V Manager. We recommend that you use a name that helps you easily identify this virtual machine, such as the name of the guest operating system or workload.

Name:

You can create a folder or use an existing folder to store the virtual machine. If you don't select a folder, the virtual machine is stored in the default folder configured for this server.

☐ Store the virtual machine in a different location

Location:

 If you plan to take snapshots of this virtual machine, select a location that has enough free space. Snapshots include virtual machine data and may require a large amount of space.

< Previous Next > Finish Cancel

5. Enter a name for this virtual machine. The name is displayed in the Hyper-V Manager.

Select **Next** to continue. The **Assign Memory** page is displayed.

New Virtual Machine Wizard

Assign Memory

Before You Begin
Specify Name and Location
Assign Memory
Configure Networking
Connect Virtual Hard Disk
Installation Options
Summary

Specify the amount of memory to allocate to this virtual machine. You can specify an amount from 8 MB through 14100 MB. To improve performance, specify more than the minimum amount recommended for the operating system.

Startup memory: MB

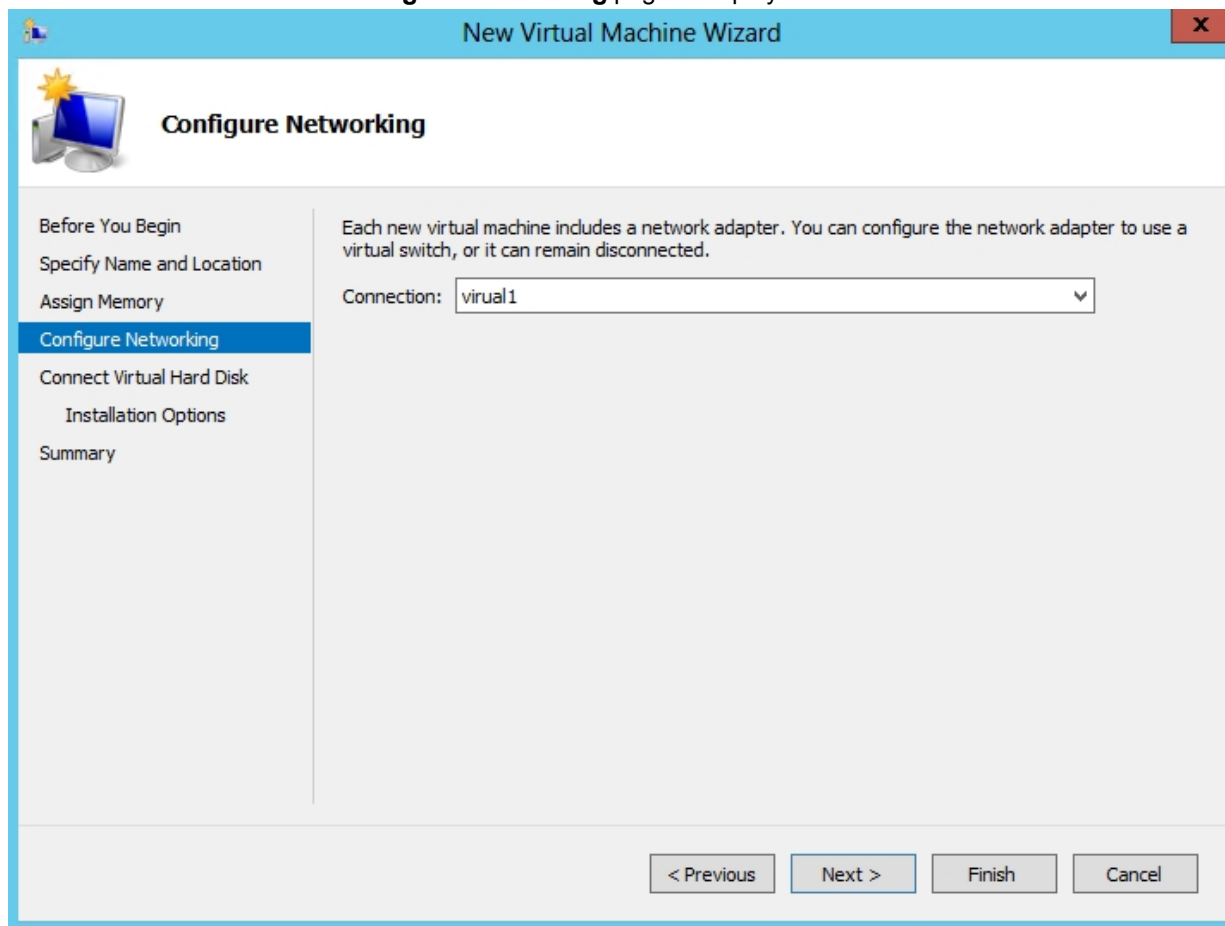
☐ Use Dynamic Memory for this virtual machine.

i When you decide how much memory to assign to a virtual machine, consider how you intend to use the virtual machine and the operating system that it will run.

< Previous Next > Finish Cancel

6. Specify the amount of memory to allocate to this virtual machine. The default memory for FortiGate VM is 1GB (1024MB).

Select **Next** to continue. The **Configure Networking** page is displayed.



7. Each new virtual machine includes a network adapter. You can configure the network adapter to use a virtual switch, or it can remain disconnected. FortiGate VM requires four network adapters. You must configure network adapters in the **Settings** page.

Select **Next** to continue. The **Connect Virtual Hard Disk** page is displayed.

New Virtual Machine Wizard

Connect Virtual Hard Disk

Before You Begin
Specify Name and Location
Assign Memory
Configure Networking
Connect Virtual Hard Disk
Summary

A virtual machine requires storage so that you can install an operating system. You can specify the storage now or configure it later by modifying the virtual machine's properties.

☐ Create a virtual hard disk
Use this option to create a dynamically expanding virtual hard disk with the default format (VHDX).

Name:
Location:
Size: GB (Maximum: 64 TB)

☒ Use an existing virtual hard disk
Use this option to attach an existing virtual hard disk, either VHD or VHDX format.

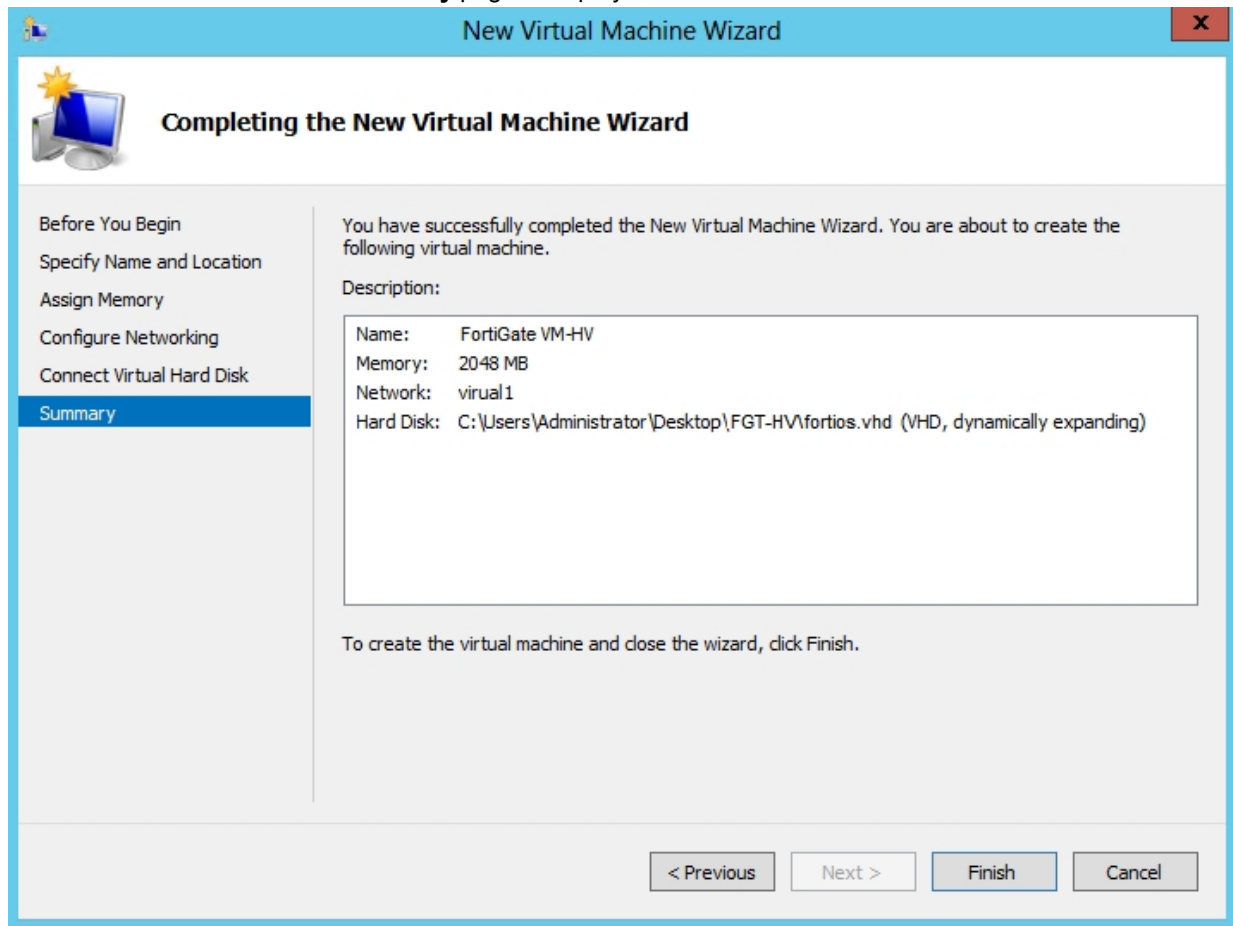
Location:

☐ Attach a virtual hard disk later
Use this option to skip this step now and attach an existing virtual hard disk later.

< Previous Next > Finish Cancel

8. Select to use an existing virtual hard disk and browse for the `fortios.vhd` file that you downloaded from the [Fortinet Customer Service & Support](#) portal.

Select **Next** to continue. The **Summary** page is displayed.



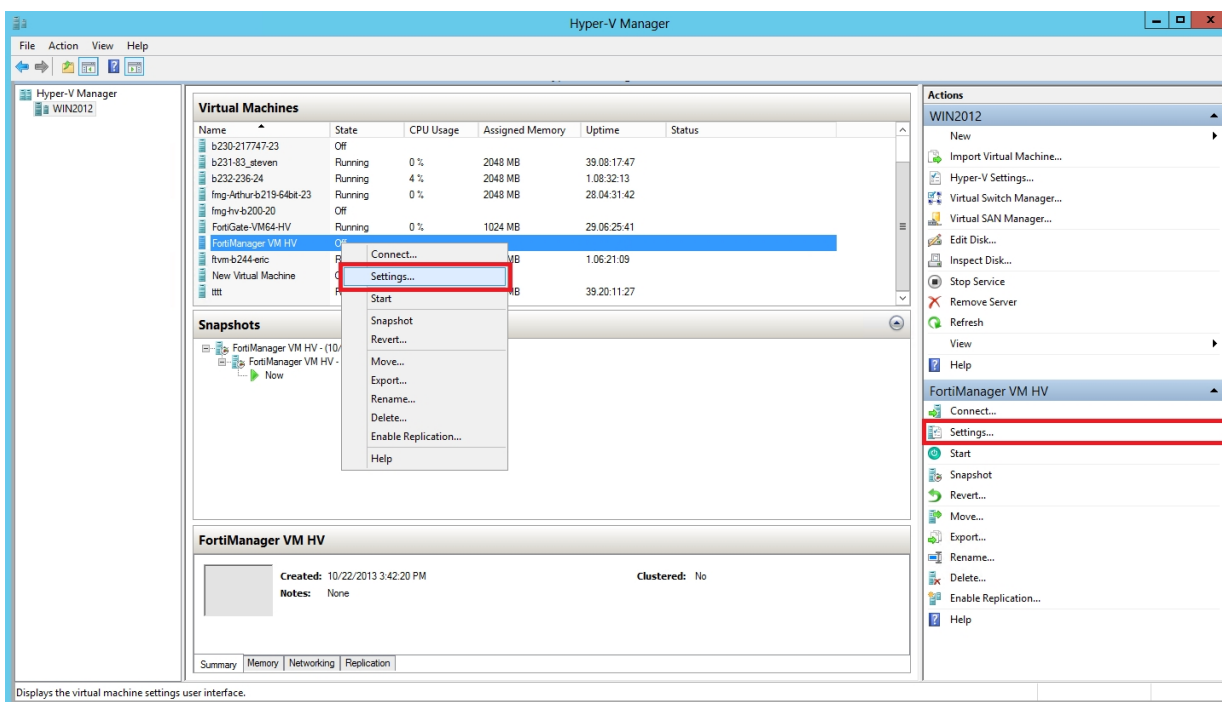
9. To create the virtual machine and close the wizard, select **Finish**.

Configure FortiGate VM hardware settings

Before powering on your FortiGate VM you must configure the virtual memory, virtual CPU, and virtual disk configuration to match your FortiGate VM license.

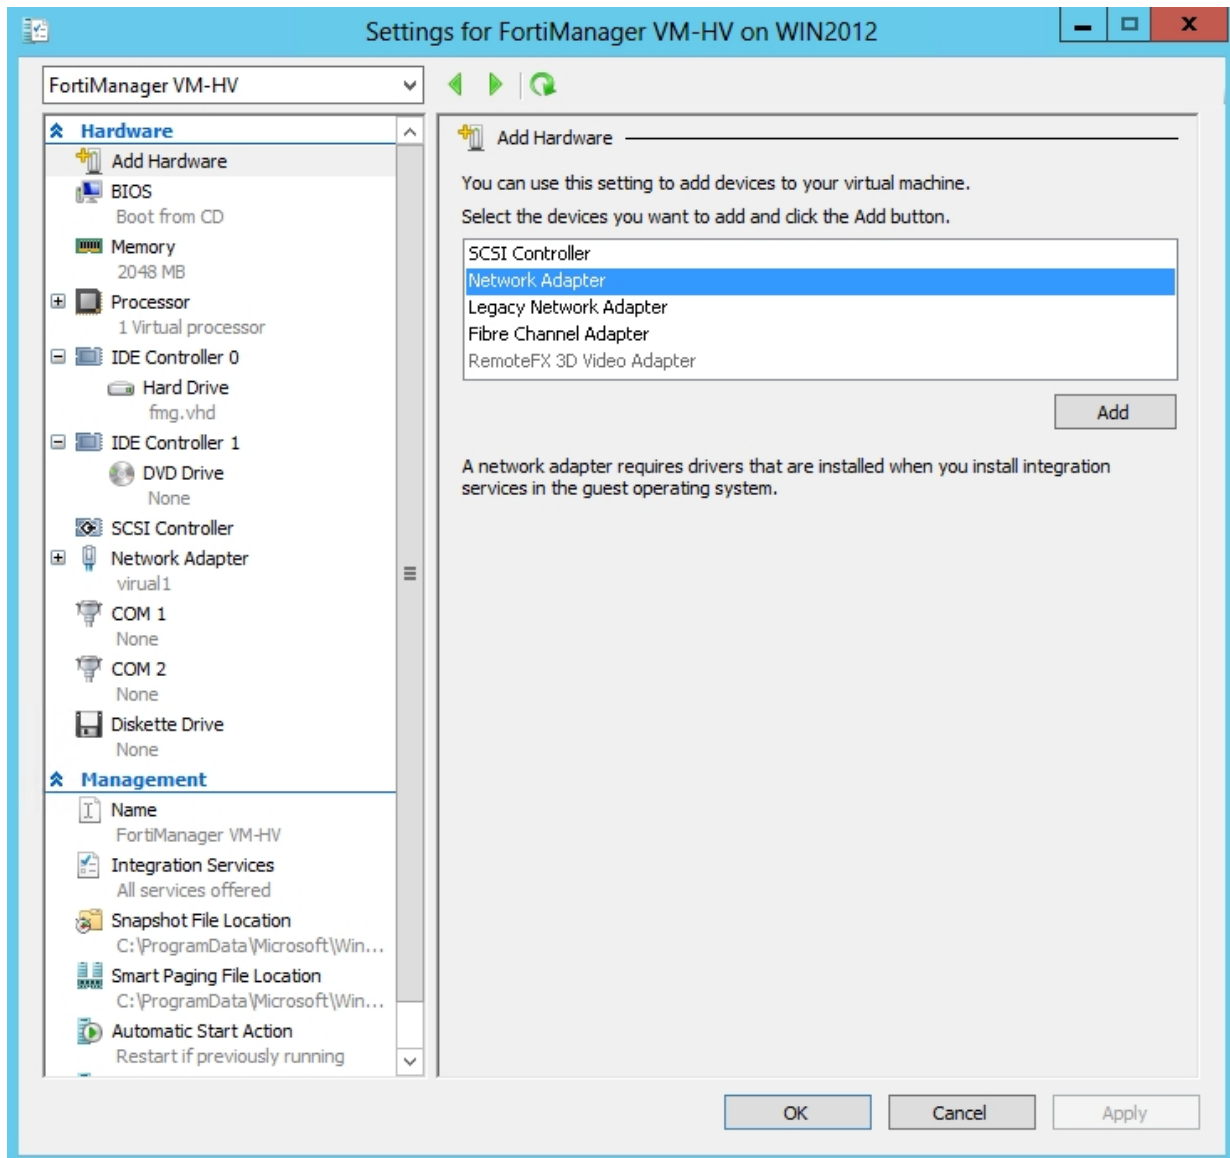
To configure settings for FortiGate VM on the server:

1. In the Hyper-V Manager, locate the name of the virtual machine, right-click the entry, and select **Settings** from the menu. Optionally, you can select the virtual machine and select **Settings** in the **Actions** menu.



Displays the virtual machine settings user interface.

The **Settings** page is displayed.



2. Configure virtual processors, network adapters, and virtual hard drive settings.
3. Select **Apply** to save the settings and then select **OK** to close the settings page.

FortiGate VM virtual processors

You must configure FortiGate VM virtual processors in the server settings page. The number of processors is dependent on your server environment.

Configure FortiGate VM virtual processors:

1. In the **Settings** page, select **Processor** from the **Hardware** menu.

The **Processor** page is displayed.

Processor

You can modify the number of virtual processors based on the number of processors on the physical computer. You can also modify other resource control settings.

Number of virtual processors:

Resource control

You can use resource controls to balance resources among virtual machines.

Virtual machine reserve (percentage):

Percent of total system resources:

Virtual machine limit (percentage):

Percent of total system resources:

Relative weight:

2. Configure the number of virtual processors for the FortiGate VM virtual machine. Optionally, you can use resource controls to balance resources among virtual machines.
3. Select **Apply** to save the settings.

FortiGate VM network adapters

You must configure FortiGate VM network adapters in the server settings page. FortiGate VM supports four network adapters.

Configure FortiGate VM network adapters:

1. In the **Settings** page, select **Add Hardware** from the **Hardware** menu, select **Network Adapter** in the device list, and select the **Add** button.

The **Network Adapter** page is displayed.

The screenshot shows the 'Network Adapter' configuration page. At the top, it says 'Specify the configuration of the network adapter or remove the network adapter.' Below this, the 'Virtual switch:' is set to 'Broadcom NetXtreme Gigabit Ethernet - Virtual Switch'. The 'VLAN ID' section has an unchecked checkbox for 'Enable virtual LAN identification' and a text box containing the value '2'. The 'Bandwidth Management' section also has an unchecked checkbox for 'Enable bandwidth management'. Below this, 'Minimum bandwidth:' and 'Maximum bandwidth:' are both set to '0 Mbps'. A note states: 'To leave the minimum or maximum unrestricted, specify 0 as the value.' At the bottom, there is a 'Remove' button and a warning icon with text: 'Use a legacy network adapter instead of this network adapter to perform a network-based installation of the guest operating system or when integration services are not installed in the guest operating system.'

1. You must manually configure four network adapters for FortiGate VM in the settings page. For each network adapter, select the virtual switch from the drop-down list.
2. Select **Apply** to save the settings.

FortiGate VM virtual hard disk


You must configure the FortiGate VM virtual hard disk in the server settings page.

If you know your environment will expand in the future, it is recommended to increase the hard disk size beyond 30GB. The VM license limit is 2TB.

Configure a FortiGate VM virtual hard drive:

1. In the **Settings** page, select **IDE Controller 0 > Hard Drive** from the **Hardware** menu.

The **Hard Drive** page is displayed.

 **Hard Drive**

You can change how this virtual hard disk is attached to the virtual machine. If an operating system is installed on this disk, changing the attachment might prevent the virtual machine from starting.


Controller: Location:

Media

You can compact or convert a virtual hard disk by editing the associated file. Specify the full path to the file.

☒ Virtual hard disk:

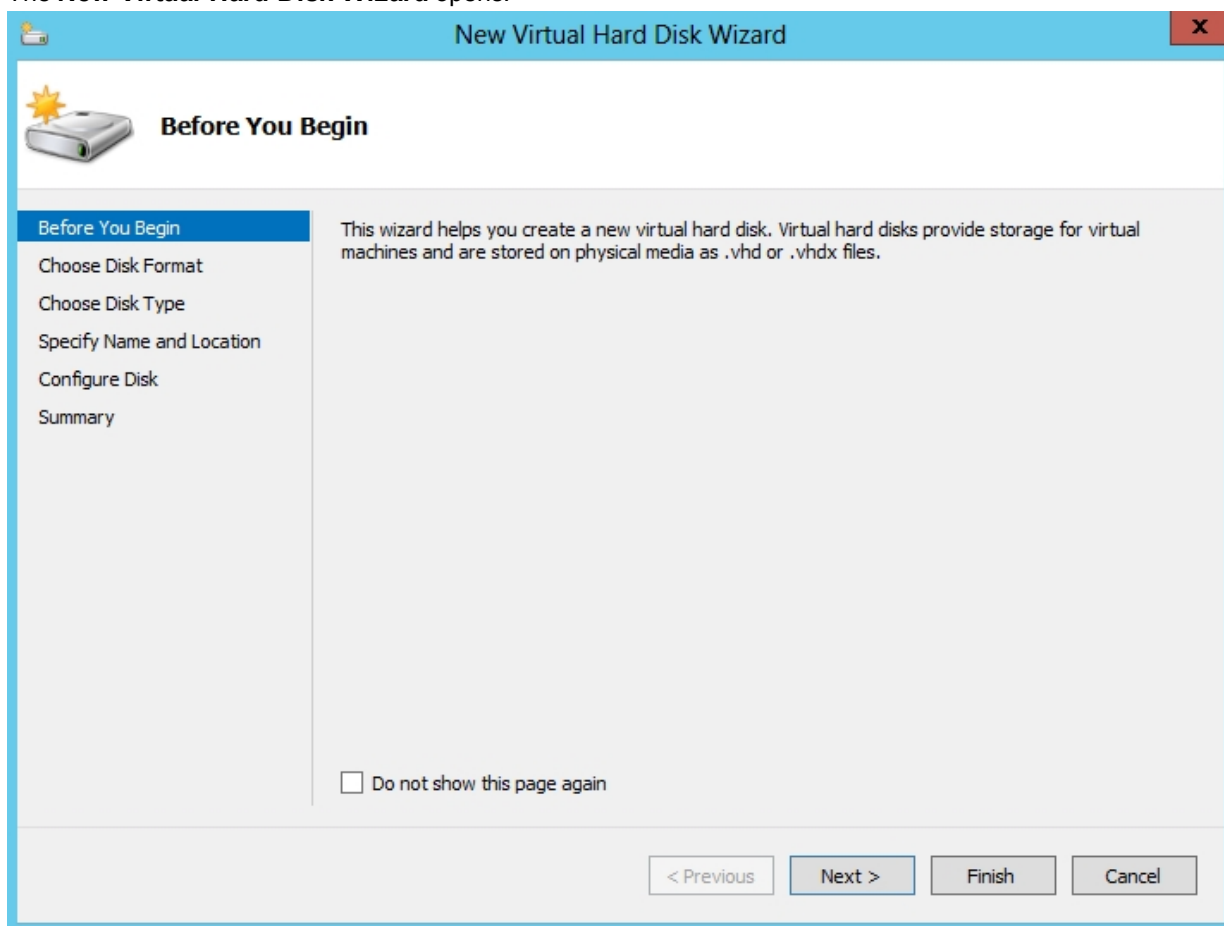
☐ Physical hard disk:

 If the physical hard disk you want to use is not listed, make sure that the disk is offline. Use Disk Management on the physical computer to manage physical hard disks.

To remove the virtual hard disk, click Remove. This disconnects the disk but does not delete the associated file.

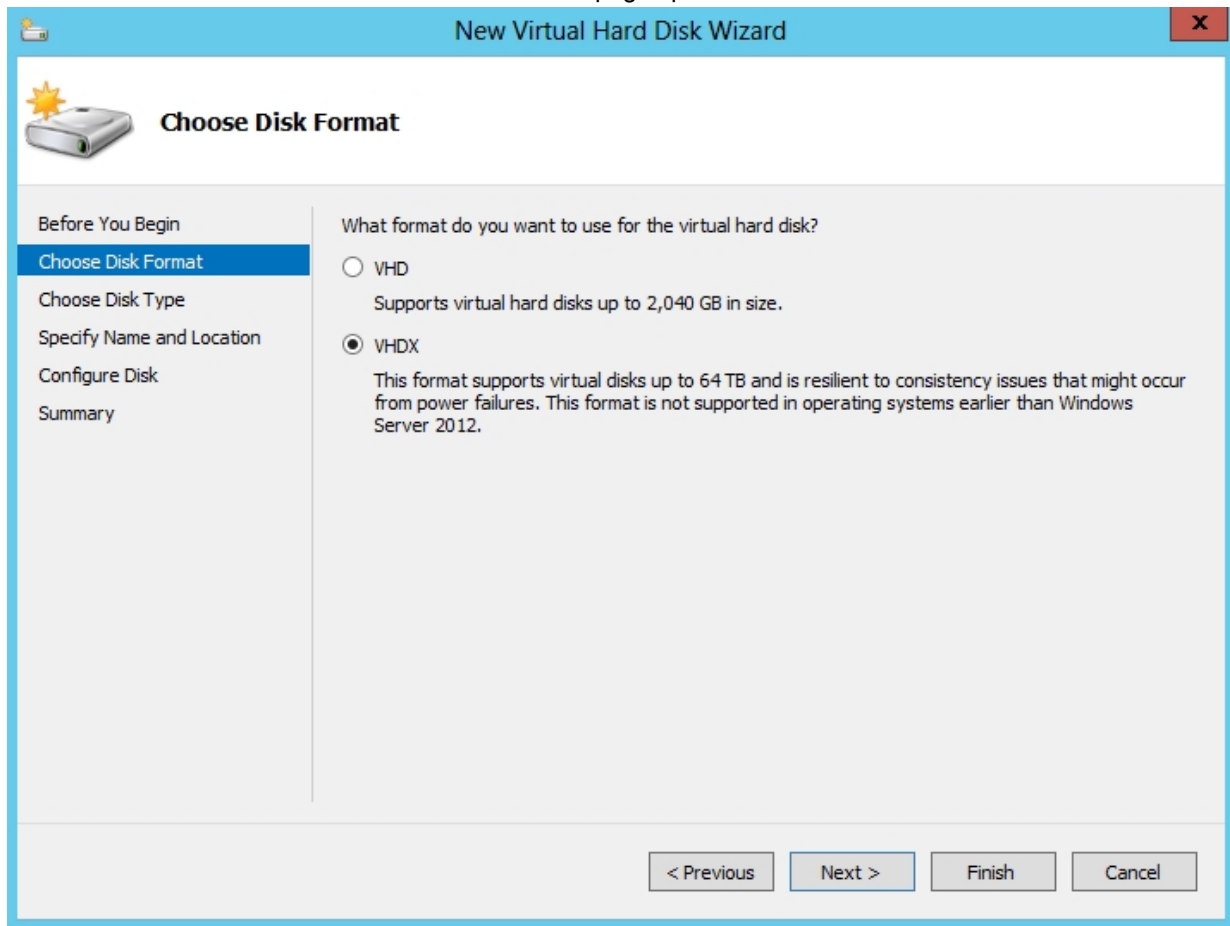
2. Select **New** to create a new virtual hard disk.

The **New Virtual Hard Disk Wizard** opens.



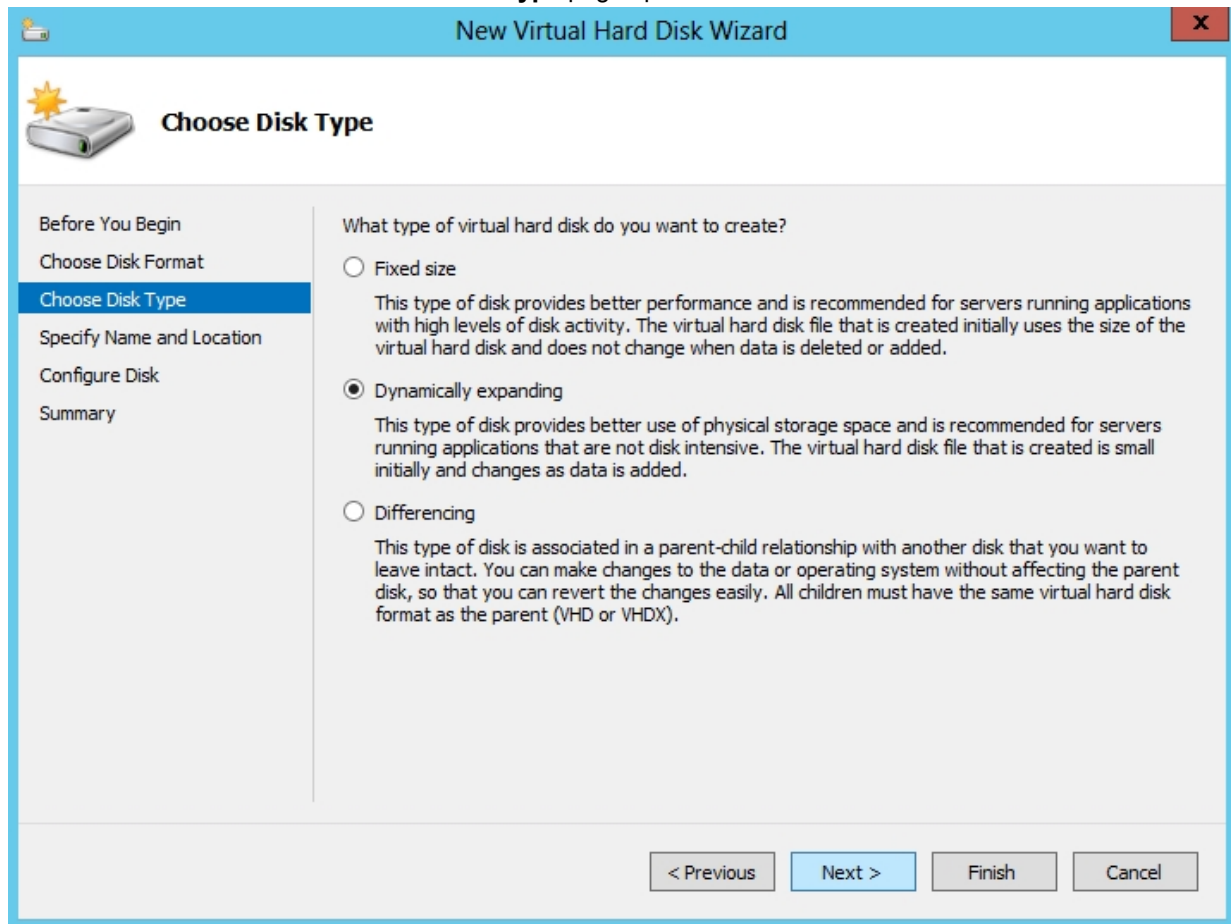
3. This wizard helps you to create a new virtual hard disk.

Select **Next** to continue. The **Choose Disk Format** page opens.



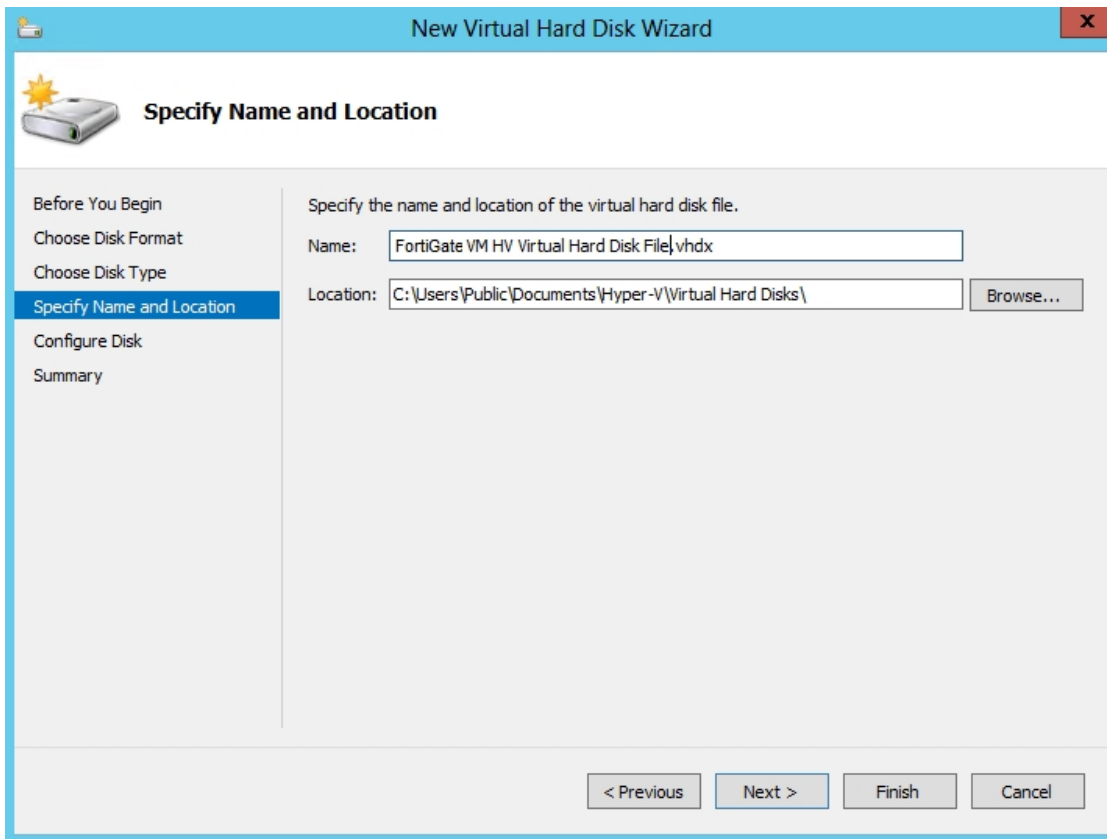
4. Select to use VHDX format virtual hard disks. This format supports virtual disks up to 64TB and is resilient to consistency issues that might occur from power failures. This format is not supported in operating systems earlier than Windows Server 2012. Note that FortiGate-VM does not support hard disks larger than 2TB.

Select **Next** to continue. The **Choose Disk Type** page opens.



5. Select the type of virtual disk you want to use. Select one of the following disk types:
- **Fixed size:** This type of disk provides better performance and is recommended for servers running applications with high levels of disk activity. The virtual hard disk file that is created initially uses the size of the virtual hard disk and does not change when data is deleted or added.
 - **Dynamic expanding:** This type of disk provides better use of physical storage space and is recommended for servers running applications that are not disk intensive. The virtual disk file that is created is small initially and changes as data is added.
 - **Differencing:** This type of disk is associated in a parent-child relationship with another disk that you want to leave intact. You can make changes to the data or operating system without affecting the parent disk, so that you can revert the changes easily. All children must have the same virtual hard disk format as the parent (VHD or VHDX).

Select **Next** to continue. The **Specify Name and Location** page opens.



6. Specify the name and location of the virtual hard disk file. Use the **Browse** button to select a specific file folder on your server.

Select **Next** to continue. The **Configure Disk** page opens.

The screenshot shows the 'New Virtual Hard Disk Wizard' window with the 'Configure Disk' step selected. The left sidebar contains the following steps: 'Before You Begin', 'Choose Disk Format', 'Choose Disk Type', 'Specify Name and Location', 'Configure Disk' (highlighted), and 'Summary'. The main area has a title 'Configure Disk' with a disk icon. Below the title, it says 'You can create a blank virtual hard disk or copy the contents of an existing physical disk.' There are two radio buttons: 'Create a new blank virtual hard disk' (selected) and 'Copy the contents of the specified physical disk:'. The 'Create a new blank virtual hard disk' option has a 'Size:' field with '30' entered and '(Maximum: 64 TB)' in parentheses. The 'Copy the contents of the specified physical disk:' option has a table listing physical hard disks. The table has two columns: 'Physical Hard Disk' and 'Size'. The rows are: '\\.\\PHYSICALDRIVE0' (1863 GB), '\\.\\PHYSICALDRIVE1' (125 MB), '\\.\\PHYSICALDRIVE2' (125 MB), '\\.\\PHYSICALDRIVE3' (125 MB), and '\\.\\PHYSICALDRIVE4' (125 MB). Below the table, there is a radio button for 'Copy the contents of the specified virtual hard disk' and a 'Path:' field with a 'Browse...' button. At the bottom, there are four buttons: '< Previous', 'Next >', 'Finish', and 'Cancel'.

Configure Disk

Before You Begin
Choose Disk Format
Choose Disk Type
Specify Name and Location
Configure Disk
Summary

You can create a blank virtual hard disk or copy the contents of an existing physical disk.

☒ Create a new blank virtual hard disk
Size: GB (Maximum: 64 TB)

☐ Copy the contents of the specified physical disk:


Physical Hard Disk	Size
\\.\PHYSICALDRIVE0	1863 GB
\\.\PHYSICALDRIVE1	125 MB
\\.\PHYSICALDRIVE2	125 MB
\\.\PHYSICALDRIVE3	125 MB
\\.\PHYSICALDRIVE4	125 MB

☐ Copy the contents of the specified virtual hard disk
Path:

< Previous Next > Finish Cancel

7. Select to **Create a new blank virtual hard disk** and enter the size of the disk in GB. The maximum size is dependent on your server environment.

Select **Next** to continue. The **Summary** page opens.



8. The summary page provides details of the virtual hard disk. Select **Finish** to create the virtual hard disk.
9. Select **Apply** to save the settings and select **OK** to exit the settings page.

High availability Hyper-V configuration

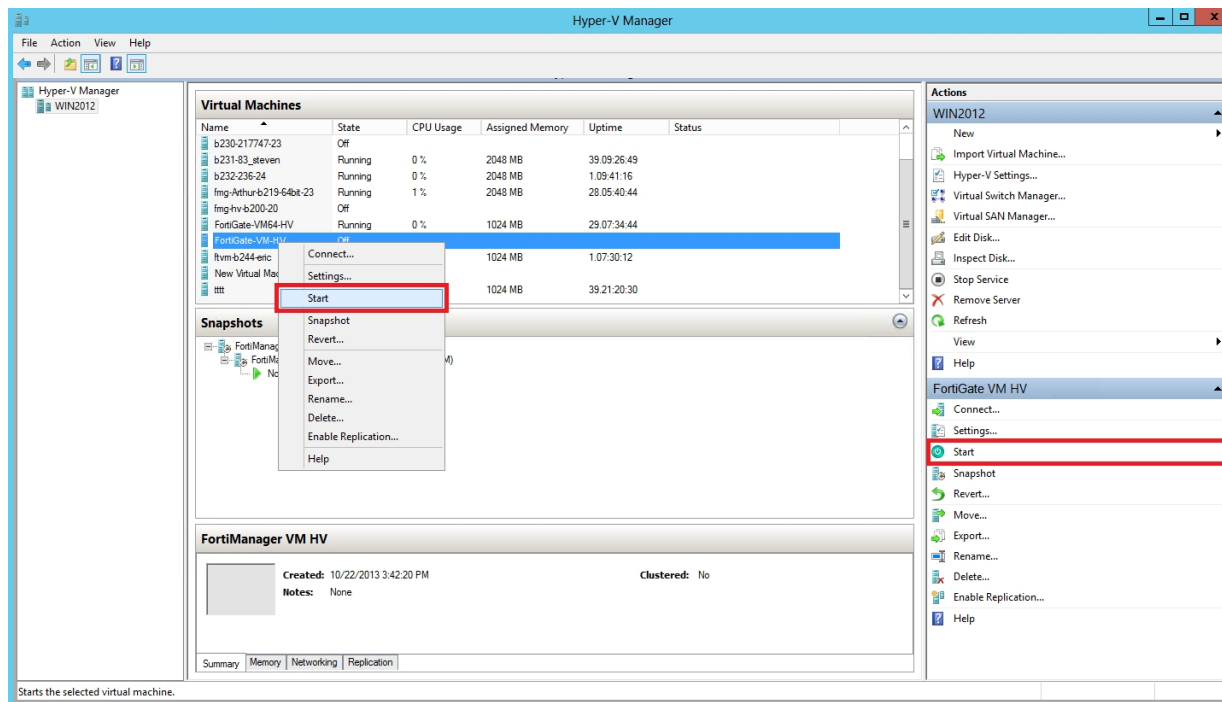
Promiscuous mode and support for MAC address spoofing is required for FortiGate-VM for Hyper-V to support FortiGate Clustering Protocol (FGCP) high availability (HA). By default the FortiGate-VM for Hyper-V has promiscuous mode enabled in the XML configuration file in the FortiGate-VM Hyper-V image. If you have problems with HA mode, confirm that this is still enabled.

In addition, because the FGCP applies virtual MAC addresses to FortiGate data interfaces and because these virtual MAC addresses mean that matching interfaces of different FortiGate-VM instances will have the same virtual MAC addresses you have to configure Hyper-V to allow MAC spoofing. But you should only enable MAC spoofing for FortiGate-VM data interfaces. You should not enable MAC spoofing for FortiGate HA heartbeat interfaces.

With promiscuous mode enabled and the correct MAC spoofing settings you should be able to configure HA between two or more FortiGate-VM for Hyper-V instances.

Start the FortiGate VM

You can now proceed to power on your FortiGate VM. Select the name of the FortiGate VM in the list of virtual machines, right-click, and select **Start** in the menu. Optionally, you can select the name of the FortiGate VM in the list of virtual machines and select **Start** in the **Actions** menu.



Deployment example – KVM

Once you have downloaded the FORTINET.out.kvm.zip file and extracted virtual hard drive image file fortios.qcow2, you can create the virtual machine in your KVM environment.

The following topics are included in this section:

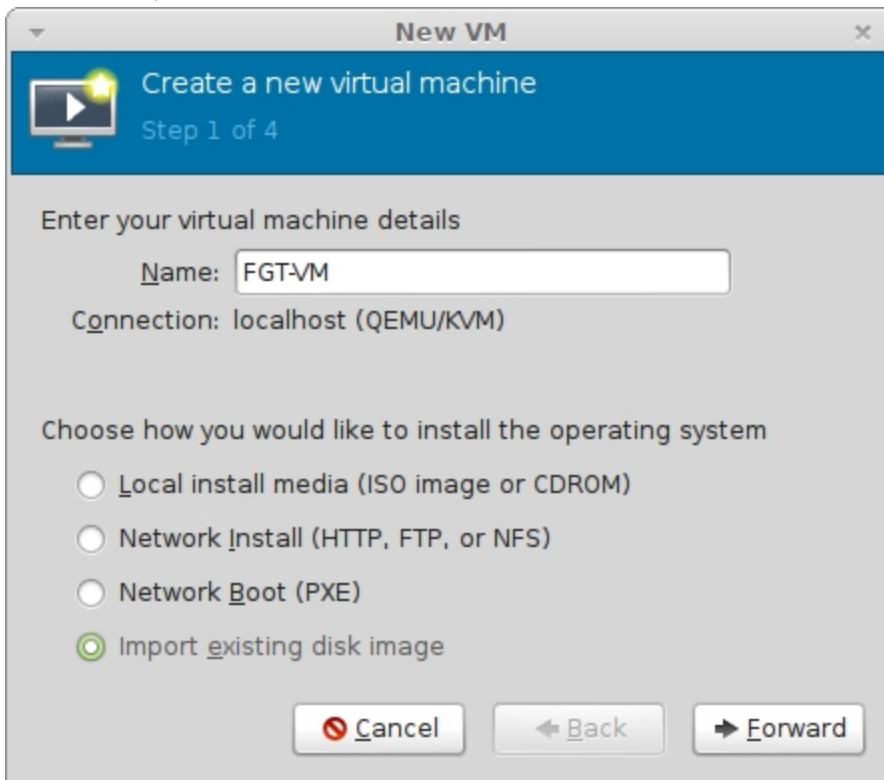
Create the FortiGate VM virtual machine

To create the FortiGate VM virtual machine:

1. Launch Virtual Machine Manager (virt-manager) on your KVM host server.

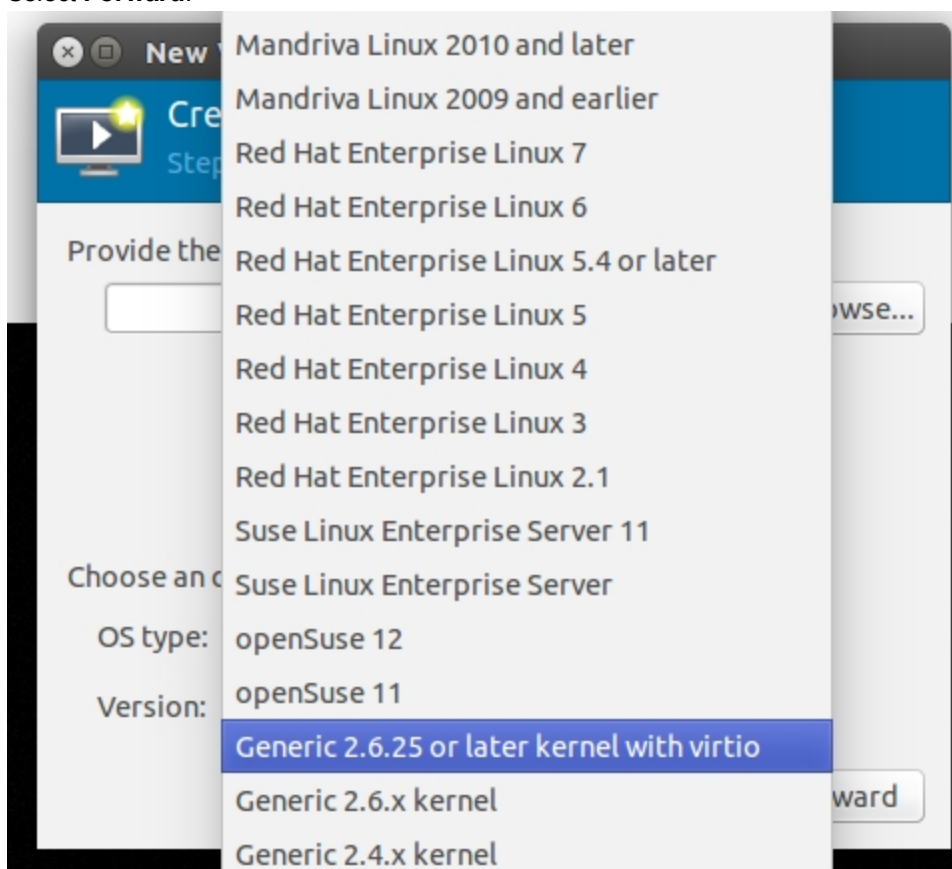
The **Virtual Machine Manager** home page opens.

2. In the toolbar, select **Create a new virtual machine**.




3. Enter a **Name** for the VM, FGT-VM for example.
4. Ensure that **Connection** is localhost. (This is the default.)
5. Select **Import existing disk image**.

6. Select **Forward**.



7. In **OS Type** select **Linux**.

8. In **Version**, select a Generic version with virtio.

9. Select **Browse**.

10. If you copied the fortios.qcow2 file to `/var/lib/libvirt/images`, it will be visible on the right. If you saved it somewhere else on your server, select **Browse Local** and find it.
11. Choose **Choose Volume**.
12. Select **Forward**.
13. Specify the amount of memory and number of CPUs to allocate to this virtual machine. The amounts must not exceed your license limits. See [Virtual FortiOS overview on page 3021](#).
14. Select **Forward**.
15. Expand **Advanced options**. A new virtual machine includes one network adapter by default. Select a network adapter on the host computer. Optionally, set a specific MAC address for the virtual network interface. Set **Virt Type** to **virtio** and **Architecture** to **qcow2**.
16. Select **Finish**.

Configure FortiGate VM hardware settings

Before powering on your FortiGate VM you must add the log disk and configure the virtual hardware of your FortiGate VM.

To configure settings for FortiGate VM on the server:

1. In the Virtual Machine Manager, locate the name of the virtual machine and then select **Open** from the toolbar.
2. Select **Add Hardware**. In the **Add Hardware** window select **Storage**.
3. Select **Create a disk image on the computer's hard drive** and set the size to 30GB.



If you know your environment will expand in the future, it is recommended to increase the hard disk size beyond 30GB. The VM license limit is 2TB.

4. Enter:

Device type	Virtio disk
Cache mode	Default
Storage format	raw



Even though raw is the storage format listed, the qcow2 format is also supported.

5. Select **Network** to configure add more the network interfaces. The **Device type** must be **Virtio**.

A new virtual machine includes one network adapter by default. You can add more through the Add Hardware window. FortiGate VM requires four network adapters. You can configure network adapters to connect to a virtual switch or to network adapters on the host computer.

6. Select **Finish**.

Start the FortiGate VM

You can now proceed to power on your FortiGate VM. Select the name of the FortiGate VM in the list of virtual machines. In the toolbar, select **Console** and then select **Start**.

Deployment example – Open Xen

Once you have downloaded the FORTINET.out.Open Xen.zip file and extracted virtual hard drive image file fortios.qcow2, you can create the virtual machine in your OpenXen environment.

The following topics are included in this section:

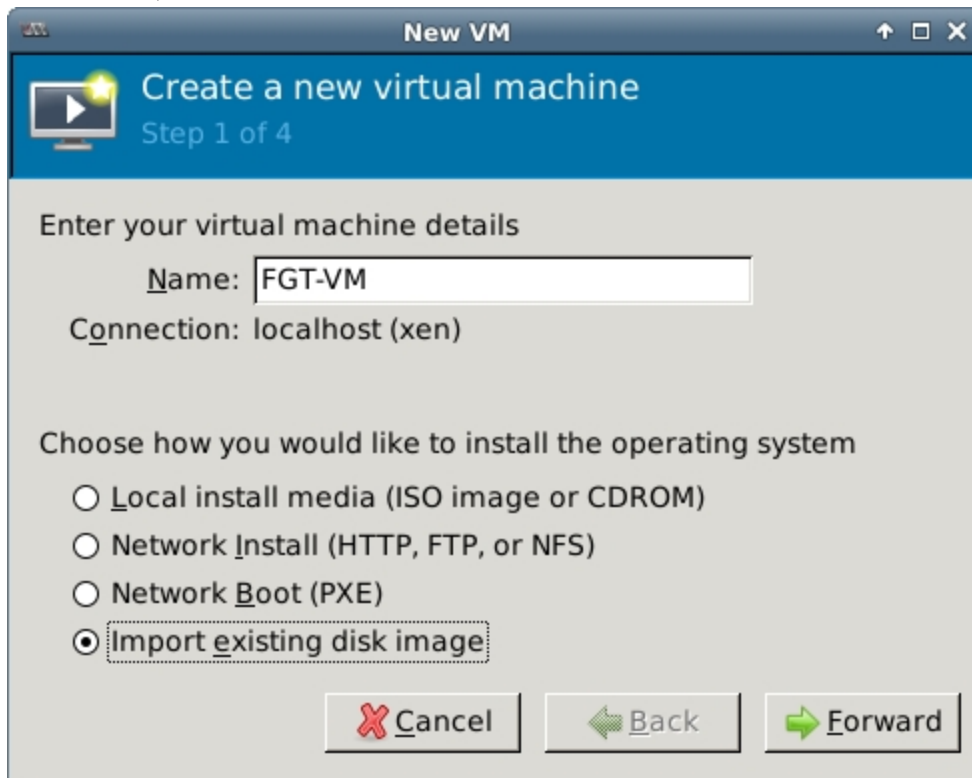
Create the FortiGate VM virtual machine

To create the FortiGate VM virtual machine:

1. Launch Virtual Machine Manager (virt-manager) on your Open Xen host server.

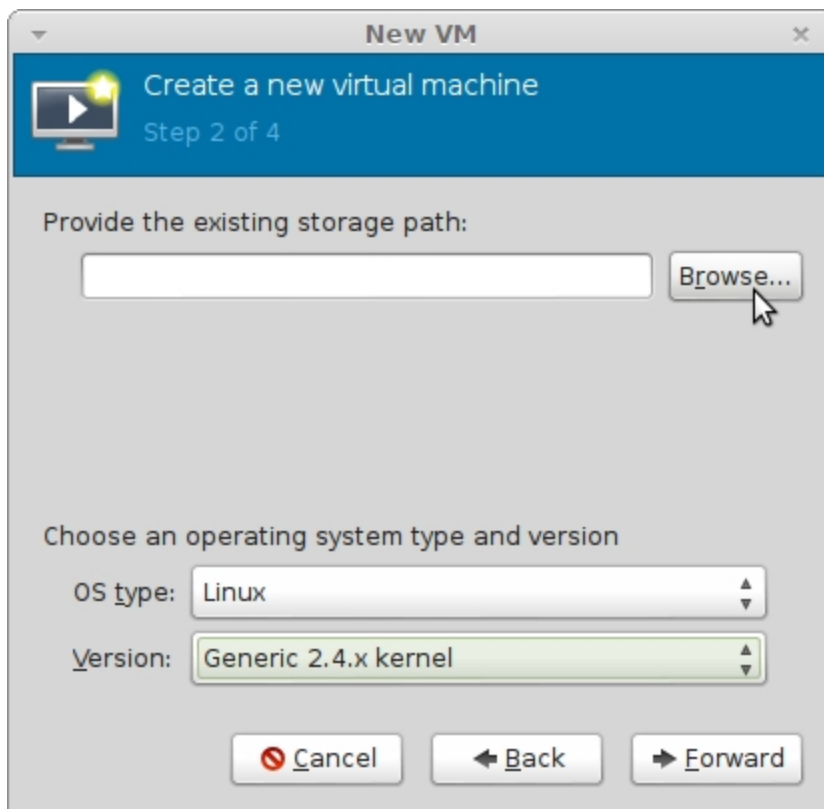
The **Virtual Machine Manager** home page opens.

2. In the toolbar, select **Create a new virtual machine**.



3. Enter a **Name** for the VM, FGT-VM for example.
4. Ensure that **Connection** is localhost. (This is the default.)
5. Select **Import existing disk image**.

6. Select **Forward**.



7. In **OS Type** select **Linux**.

8. In **Version**, select **Generic 2.4.x.kernel**.

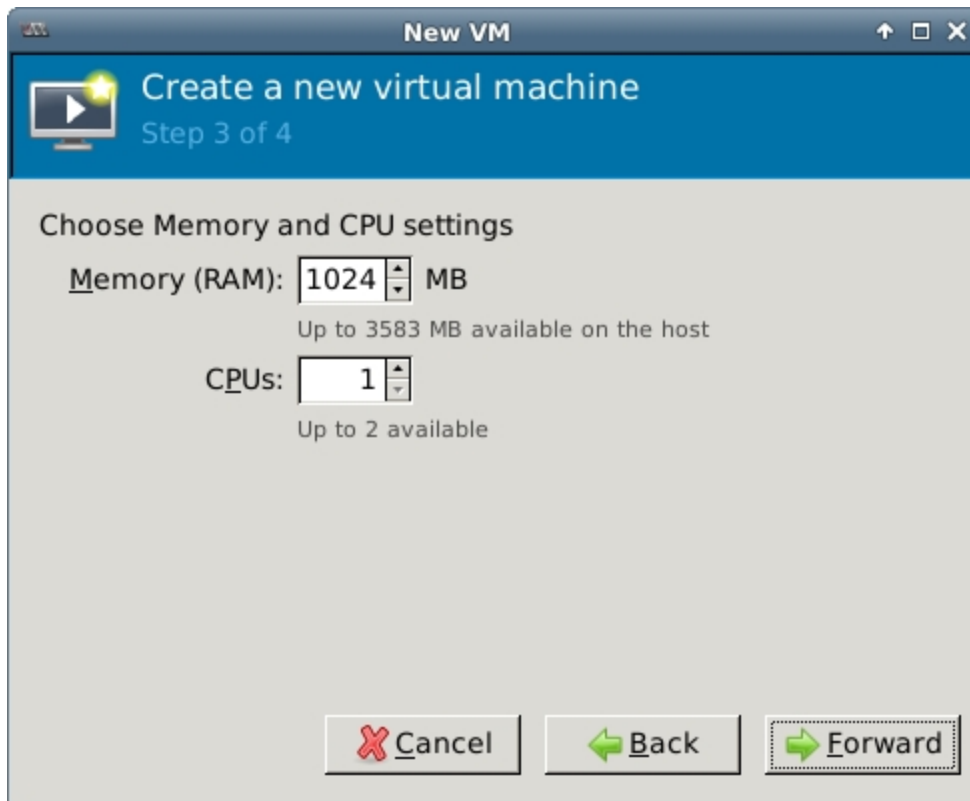
9. Select **Browse**.

The **Locate or create storage volume** window opens.

10. Select **Browse Local**, find the fortios.qcow2 disk image file.

11. Select fortios.qcow2 and select **Choose Volume**.

12. Select **Forward**.



New VM

Create a new virtual machine
Step 3 of 4

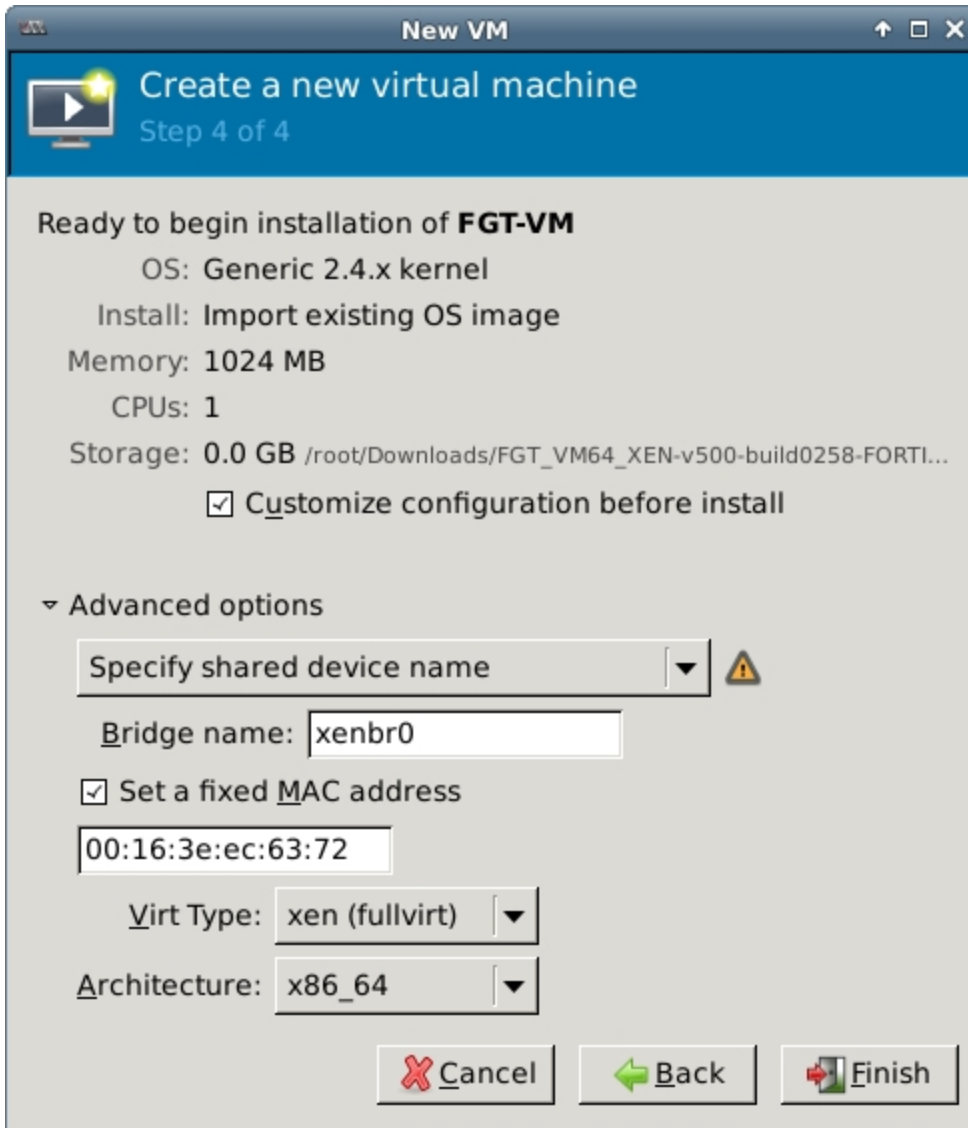
Choose Memory and CPU settings

Memory (RAM): 1024 MB
Up to 3583 MB available on the host

CPUs: 1
Up to 2 available

Cancel Back Forward

13. Specify the amount of memory and number of CPUs to allocate to this virtual machine. The amounts must not exceed your license limits.

14. Select Forward.

New VM

Create a new virtual machine
Step 4 of 4

Ready to begin installation of **FGT-VM**

OS: Generic 2.4.x kernel

Install: Import existing OS image

Memory: 1024 MB

CPUs: 1

Storage: 0.0 GB /root/Downloads/FGT_VM64_XEN-v500-build0258-FORTI...

☒ **Customize configuration before install**

▼ **Advanced options**

Specify shared device name ▼ ⚠

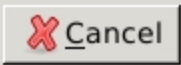
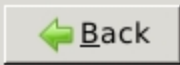
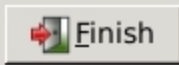
Bridge name: xenbr0

☒ Set a fixed MAC address

00:16:3e:ec:63:72

Virt Type: xen (fullvirt) ▼

Architecture: x86_64 ▼

- 15. Select **Customize configuration before install**.** This enables you to make some hardware configuration changes before VM creation is started.
- 16. Expand **Advanced options**.** A new virtual machine includes one network adapter by default. Select **Specify shared device name** and enter the name of the bridge interface on the Open Xen host. Optionally, set a specific MAC address for the virtual network interface. **Virt Type** and **Architecture** are set by default and should be correct.
- 17. Select **Finish**.**

The virtual machine hardware configuration window opens.



You can use this window to add hardware such as network interfaces and disk drives.

18. Select **Add Hardware**. In the **Add Hardware** window select **Storage**.
19. Select **Create a disk image on the computer's hard drive** and set the size to 30GB.



If you know your environment will expand in the future, it is recommended to increase the hard disk size beyond 30GB. The VM license limit is 2TB.

20. Enter:

Device type	Virtio disk
Cache mode	Default
Storage format	raw

21. Select **Network** to configure add more the network interfaces. The **Device type** must be **Virtio**.

A new virtual machine includes one network adapter by default. You can add more through the Add Hardware window. FortiGate VM requires four network adapters. You can configure network adapters to connect to a virtual

switch or to network adapters on the host computer.

22. Select **Finish.**

23. Select **Begin Installation.** After the installation completes successfully, the VM starts and the console window opens.

Deployment example – Citrix XenServer

Once you have downloaded the FORTINET.out.CitrixXen.zip file and extracted the files, you can create the virtual machine in your Citrix Xen environment.

The following topics are included in this section:

Create the FortiGate VM virtual machine (XenCenter)

To create the FortiGate VM virtual machine from the OVF file

1. Launch XenCenter on your management computer.

The management computer can be any computer that can run Citrix XenCenter, a Windows application.


2. If you have not already done so, select **ADD a server**. Enter your Citrix XenServer IP address and the root logon credentials required to manage that server.

Your Citrix XenServer is added to the list in the left pane.

The **Virtual Machine Manager** home page opens.

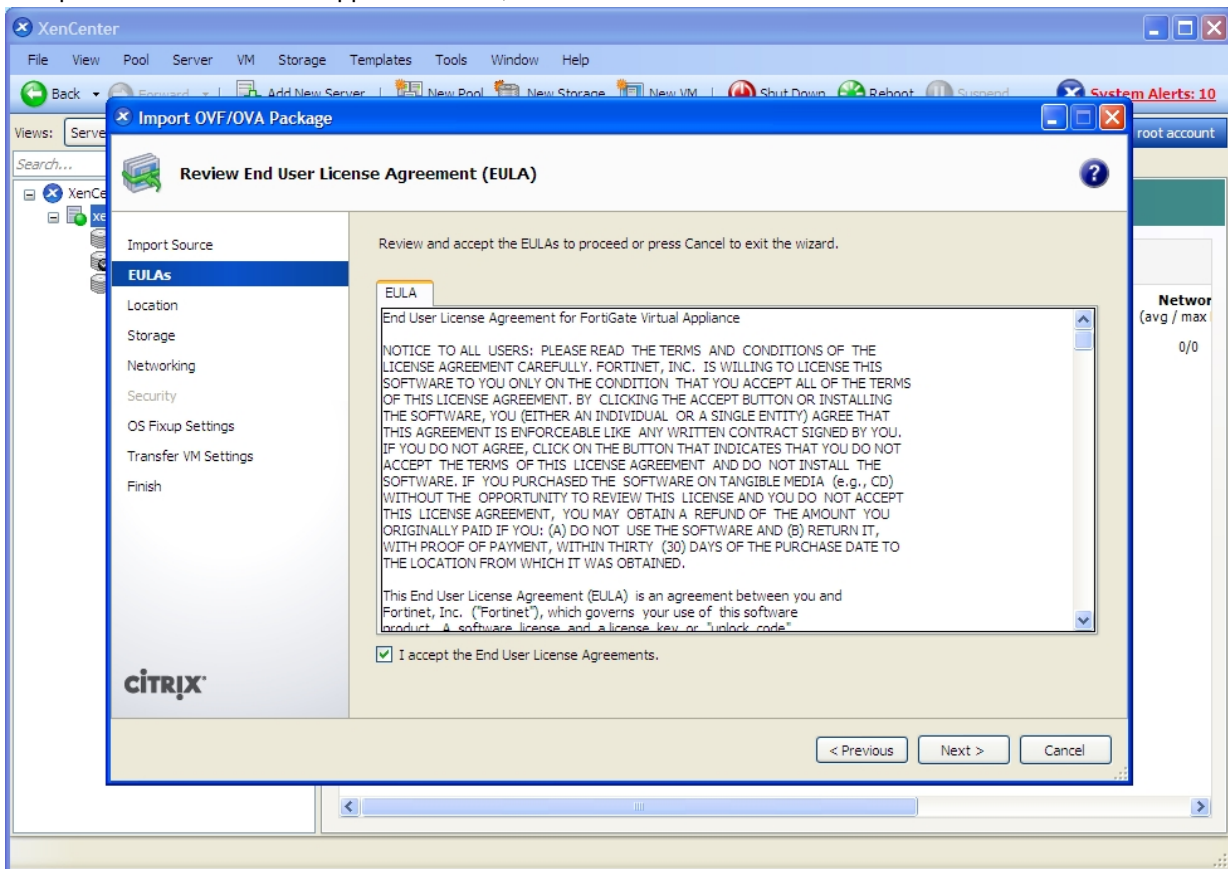
3. Go to **File > Import**. An import dialog will appear.

- Click the **Browse** button, find the FortiGate-VM64-Xen.ovf template file, then click **Open**.



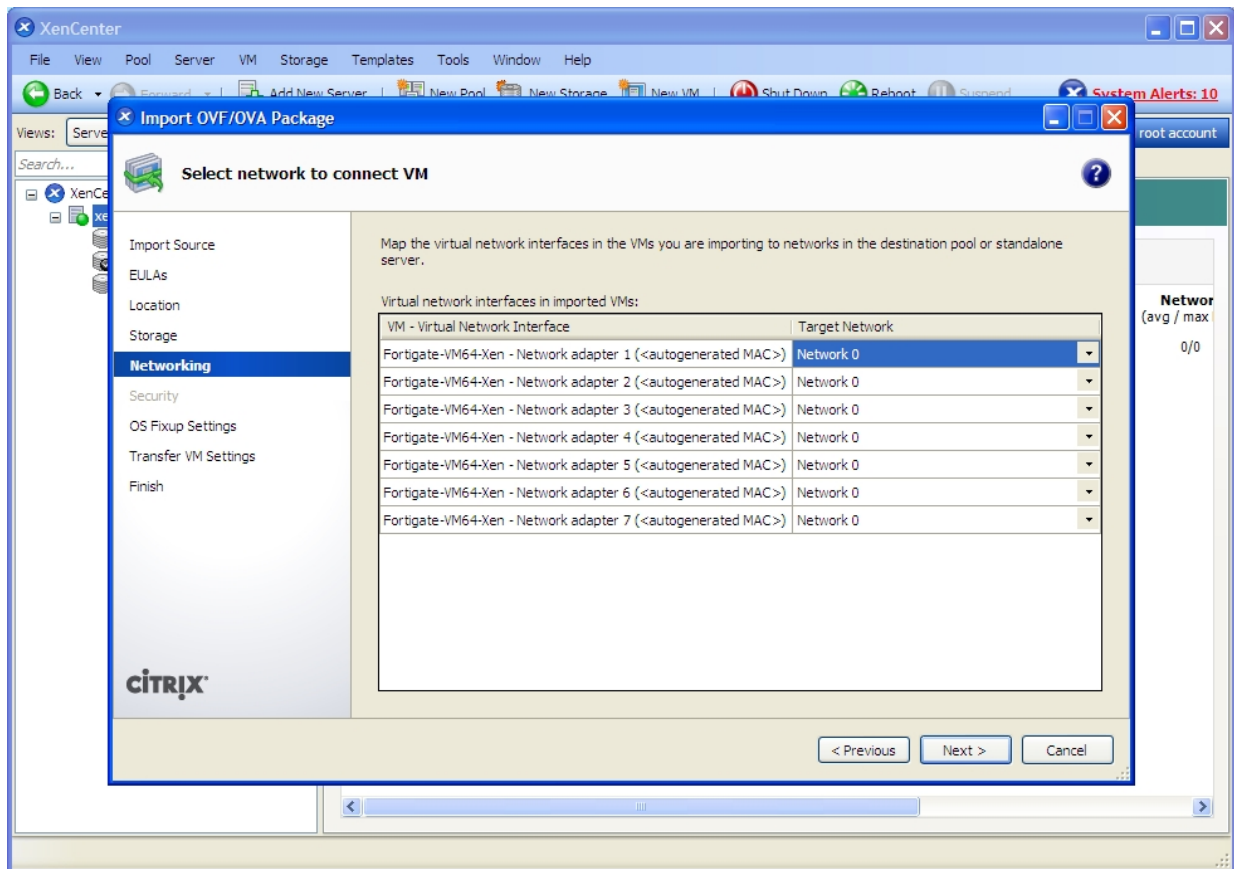
- Select **Next**.

6. Accept the FortiGate Virtual Appliance EULA, then select **Next**.



7. Choose the pool or standalone server that will host the VM, then select **Next**.
8. Select the storage location for FortiGate VM disk drives or accept the default. Select **Next**.

9. Configure how each vNIC (virtual network adapter) in FortiGate VM will be mapped to each vNetwork on the Citrix XenServer, then click **Next**.



10. Click **Next** to skip OS fixup.
11. Select **Next** to use the default network settings for transferring the VM to the host.
12. Select **Finish**.

The Citrix XenServer imports the FortiGate VM files and configures the VM as specified in the OVF template. Depending on your computer's hardware speed and resource load, and also on the file size and speed of the network connection, this might take several minutes to complete.



When VM import is complete, the XenCenter left pane includes the FortiGate VM in the list of deployed VMs for your Citrix XenServer.

Configure virtual hardware


Before you start your FortiGate-VM for the first time, you need to adjust your virtual machine's virtual hardware settings to meet your network requirements.

Configuring number of CPUs and memory size

Your FortiGate-VM license limits the number CPUs and amount of memory that you can use. The amounts you allocate must not exceed your license limits.

To access virtual machine settings

1. Open XenCenter.
2. Select your FortiGate VM in the left pane.
The tabs in the right pane provide access to the virtual hardware configuration. The Console tab provides access to the FortiGate console.
1. To set the number of CPUs
2. In the XenCenter left pane, right-click the FortiGate VM and select Properties.
The Properties window opens.
3. In the left pane, select CPU.



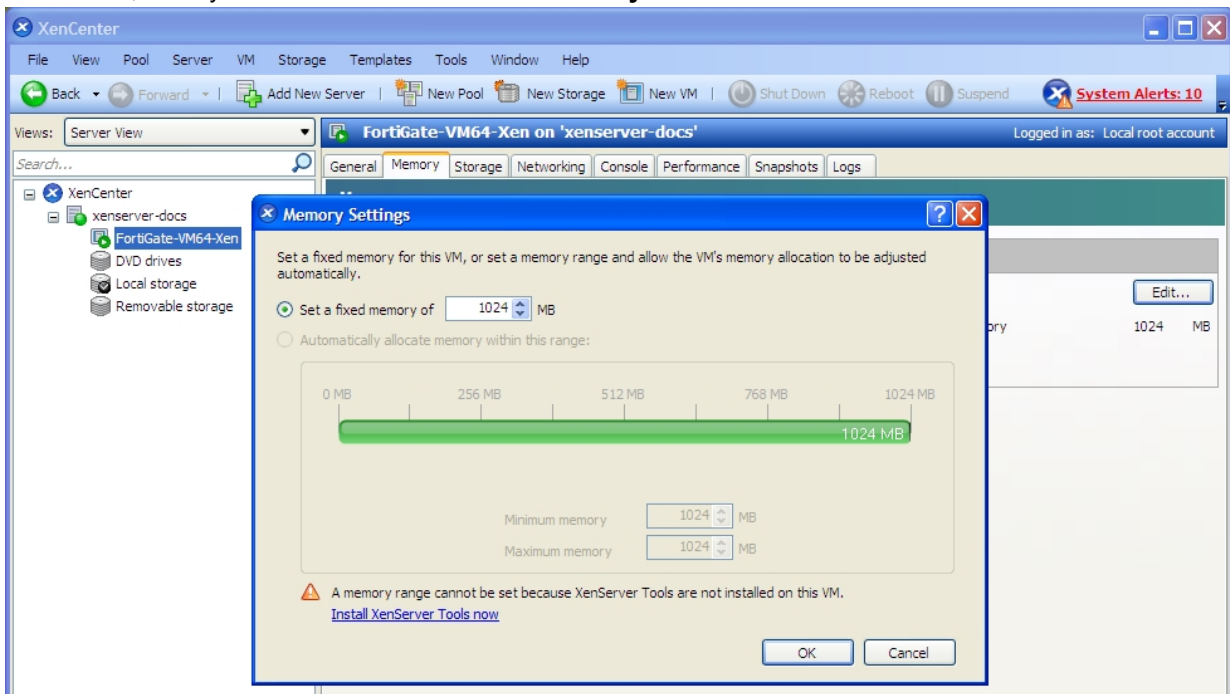
4. Adjust **Number of CPUs** and then select **OK**.

XenCenter will warn if you select more CPUs than the Xen host computer contains. Such a configuration might reduce performance.

To set memory size

1. In the XenCenter left pane, select the FortiGate VM.
2. In the right pane, select the **Memory** tab.

3. Select **Edit**, modify the value in the **Set a fixed memory of** field and select **OK**.



Configuring disk storage

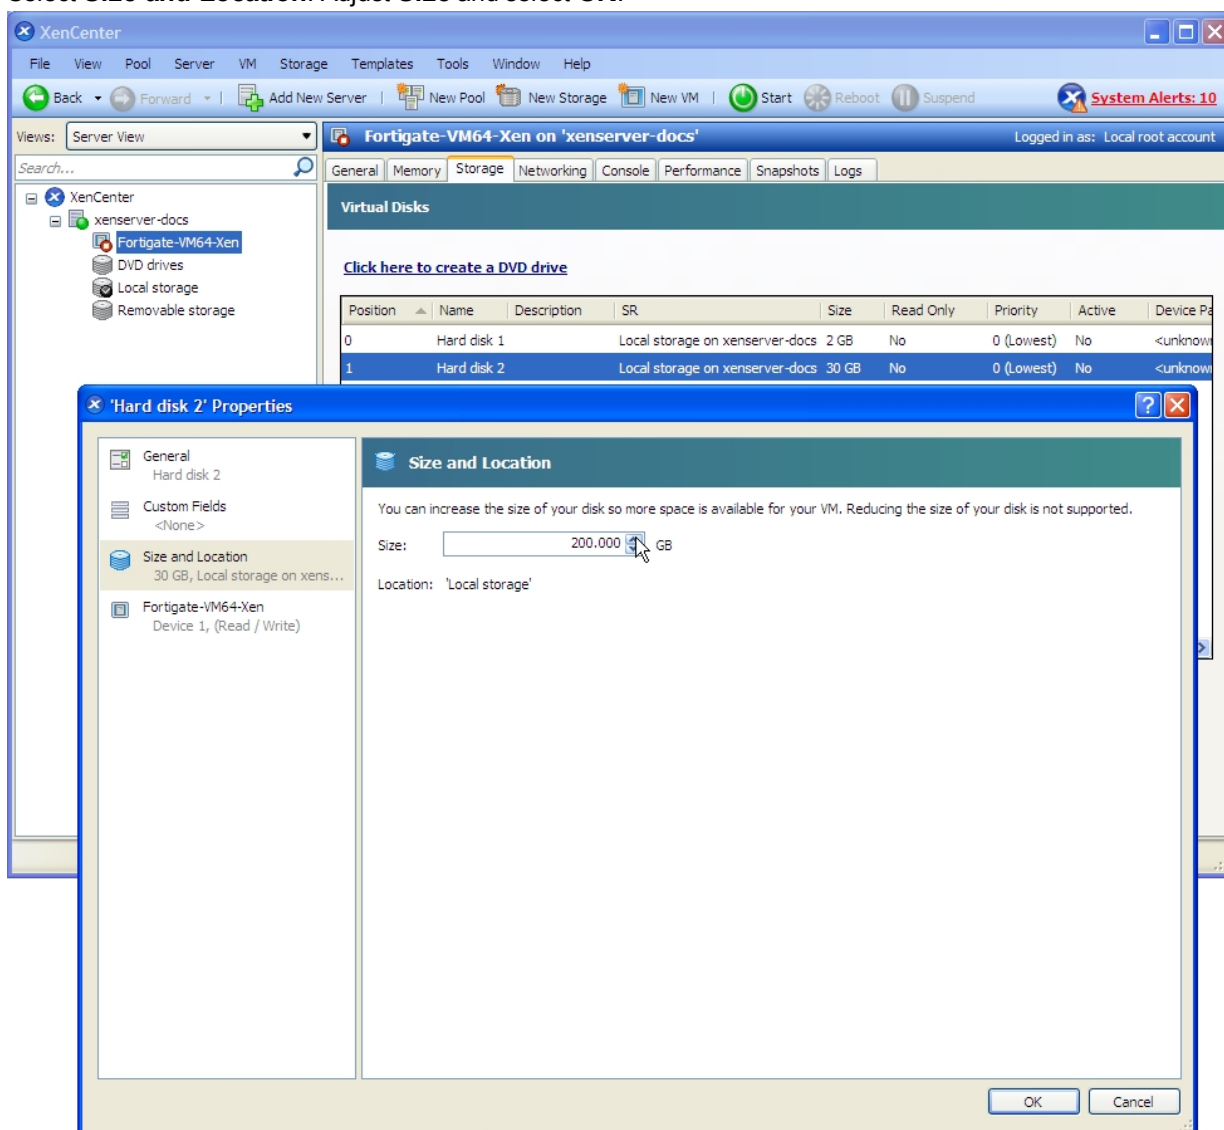
By default the FortiGate VM data disk 30GB. You will probably want to increase this. Disk resizing must be done before you start the VM for the first time.

To resize the FortiGate data disk

1. In the XenCenter left pane, select the FortiGate VM.
2. Select the **Storage** tab. Select **Hard disk 2** (the 30GB drive), then select **Properties**.

The '**Hard disk 2** Properties' window opens.

3. Select **Size and Location**. Adjust **Size** and select **OK**.



FortiGate VM initial configuration

Before you can connect to the FortiGate VM web-based manager you must configure a network interface in the FortiGate VM console. Once an interface with administrative access is configured, you can connect to the FortiGate VM web-based Manager and upload the FortiGate VM license file that you downloaded from the [Customer Service & Support](#) website.

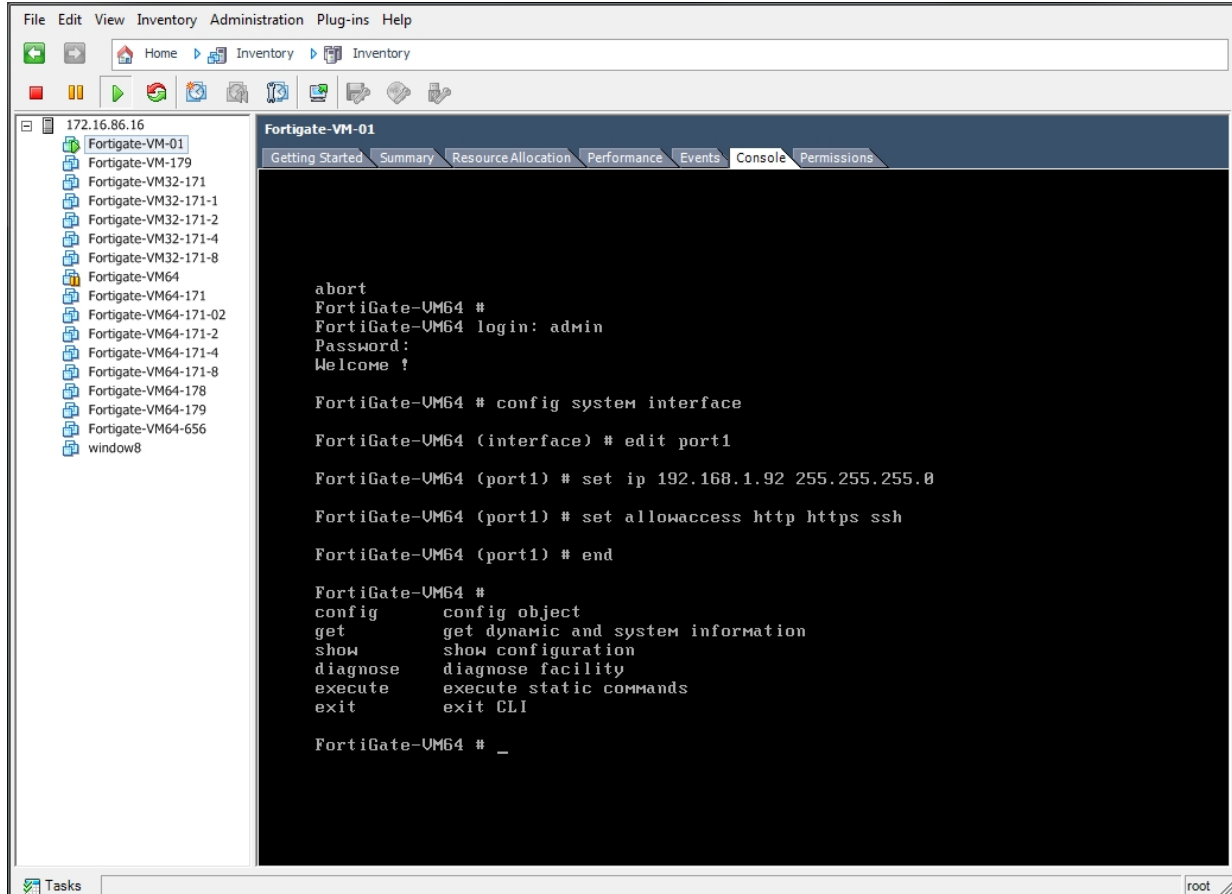
The following topics are included in this section:

Set FortiGate VM port1 IP address

Hypervisor management environments include a guest console window. On the FortiGate VM, this provides access to the FortiGate console, equivalent to the console port on a hardware FortiGate unit. Before you can access the Web-based manager, you must configure FortiGate VM port1 with an IP address and administrative access.

To configure the port1 IP address:

1. In your hypervisor manager, start the FortiGate VM and access the console window. You might need to press Return to see a login prompt.

Example of FortiGate VM console access:

2. At the FortiGate VM login prompt enter the username `admin`. By default there is no password. Just press Return.
3. Using CLI commands, configure the port1 IP address and netmask. Also, HTTP access must be enabled because until it is licensed the FortiGate VM supports only low-strength encryption. HTTPS access will not work.

For example:

```

config system interface
edit port1
set ip 192.168.0.100 255.255.255.0
append allowaccess http
end
  
```



You can also use the `append allowaccess` CLI command to enable other access protocols, such as `auto-ipsec`, `http`, `probe-response`, `radius-acct`, `snmp`, and `telnet`. The `ping`, `https`, `ssh`, and `fgfm` protocols are enabled on the `port1` interface by default.

4. To configure the default gateway, enter the following CLI commands:

```
config router static
edit 1
set device port1
set gateway <class_ip>
end
```



You must configure the default gateway with an IPv4 address. FortiGate VM needs to access the Internet to contact the FortiGuard Distribution Network (FDN) to validate its license.

5. To configure your DNS servers, enter the following CLI commands:

```
config system dns
set primary <Primary DNS server>
set secondary <Secondary DNS server>
end
```



The default DNS servers are 208.91.112.53 and 208.91.112.52.

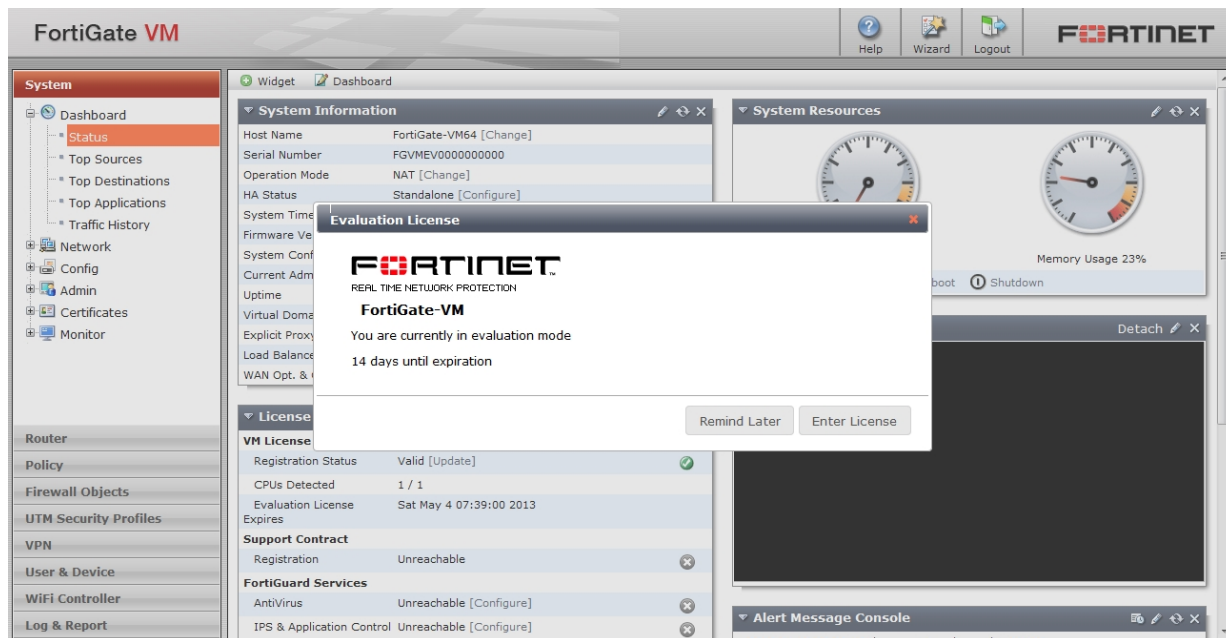
6. To upload the FortiGate VM license from an FTP or TFTP server, use the following CLI command:

```
execute restore vmlicense {ftp | tftp} <VM license file name> <Server IP or FQDN>
[:server port]
```



You can also upload the license in the FortiGate VM Web-based Manager. See [Set FortiGate VM port1 IP address on page 3082](#).

Web-based Manager and Evaluation License dialog box



Connect to the FortiGate VM web-based manager

When you have configured the port1 IP address and netmask, launch a web browser and enter the IP address that you configured for port1. At the login page, enter the username `admin` and password field and select **Login**. The default password is no password. The Web-based Manager will appear with an **Evaluation License** dialog box.



Due to low encryption on the FortiGate side of the connection, some modern browsers will not allow the connection. If that is the case, the adjusting of the browser settings is unlikely to make a difference. There are two options:

- Use FTP or TFTP to apply the license
- Make sure that the interface on the port being accessed has been configured to allow HTTP/HTTPS access

Upload the FortiGate VM license file

Every Fortinet VM includes a 15-day trial license. During this time the FortiGate VM operates in evaluation mode. Before using the FortiGate VM you must enter the license file that you downloaded from the [Customer Service & Support](#) website upon registration.

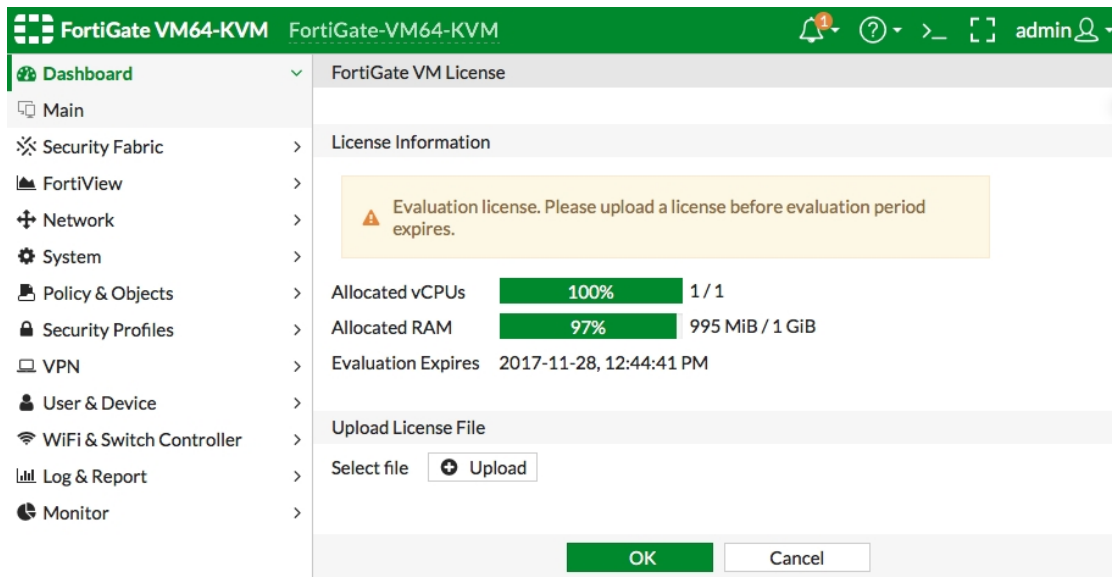
GUI

To upload the FortiGate VM license file:

1. There are 2 ways to get to the License upload window.

- i. In the **Dashboard > Main** window, in the **Virtual Machine** widget, left click on the **FGVMEV** (FortiGate-VM Evaluation) **License** icon. This will reveal a menu of selections to take you to directly to the **FortiGate VM License** window or to the **FortiGuard Details** window.
 - ii. Go to **System > FortiGuard**. In the **Larcentist Information** section, go to the **Virtual Machine** row and click on the link to **FortiGate VM License**.
2. In the **Evaluation License** dialog box, select **Enter License**.
The license upload page opens.

License upload page:



3. Select **Upload** and locate the license file (.lic) on your computer. Select **OK** to upload the license file.
4. Refresh the browser to login.
5. Enter `admin` in the Name field and select **Login**. The VM registration status appears as valid in the License Information widget once the license has been validated by the FortiGuard Distribution Network (FDN) or FortiManager for closed networks.



Modern browsers can have an issue with allowing connecting to a FortiGate if the encryption on the device is too low. Adjusting browser settings does not normally mitigate the issue. If this happens, Admins must use a FTP/TFTP server to apply the license.

CLI

You can also upload the license file via the CLI using the following CLI command:

```
execute restore vmlicense {ftp | tftp} <filename string> <ftp server>[:ftp port]
```

Example:

The following is an example output when using a tftp server to install license.

```
exec restore vmlicense tftp license.lic 10.0.1.2
This operation will overwrite the current VM license!Do you want to continue? (y/n)y
```

```
Please wait...Connect to tftp server 10.0.1.2 ...
Get VM license from tftp server OK.
VM license install succeeded.
Rebooting firewall.
```



The command has the side effect of rebooting the firewall without giving you a chance to back out or delay the reboot, so be careful about the timing of using the command.

Validate the FortiGate VM license with FortiManager

You can validate your FortiGate VM license with some models of FortiManager. To determine whether your FortiManager unit has the VM Activation feature, see Features section of the [FortiManager Product Data sheet](#).

To validate your FortiGate VM with your FortiManager:

1. To configure your FortiManager as a closed network, enter the following CLI command on your FortiManager:

```
config fmupdate publicnetwork
  set status disable
end
```
2. To configure FortiGate VM to use FortiManager as its override server, enter the following CLI commands on your FortiGate VM:

```
config system central-management
  set mode normal
  set type fortimanager
  set fmg <IPv4 address of the FortiManager device>
  set fmg-source-ip <Source IPv4 address when connecting to the FortiManager device>
  set include-default-servers disable
  set vdom <Enter the name of the VDOM to use when communicating with the FortiManager device>
end
```
3. Load the FortiGate VM license file in the Web-based Manager. Go to **System > Dashboard > Status**. In the **License Information** widget, in the **Registration Status** field, select **Update**. Browse for the `.lic` license file and select **OK**.
4. To activate the FortiGate VM license, enter the following CLI command on your FortiGate VM:

```
execute update-now
```
5. To check the FortiGate VM license status, enter the following CLI commands on your FortiGate VM:

```
get system status
```

The following output is displayed:

```
Version: Fortigate-VM v5.0,build0099,120910 (Interim)
Virus-DB: 15.00361(2011-08-24 17:17)
Extended DB: 15.00000(2011-08-24 17:09)
Extreme DB: 14.00000(2011-08-24 17:10)
IPS-DB: 3.00224(2011-10-28 16:39)
FortiClient application signature package: 1.456(2012-01-17 18:27)
Serial-Number: FGVM02Q105060000
License Status: Valid
BIOS version: 04000002
Log hard disk: Available
```

```

Hostname: Fortigate-VM
Operation Mode: NAT
Current virtual domain: root
Max number of virtual domains: 10
Virtual domains status: 1 in NAT mode, 0 in TP mode
Virtual domain configuration: disable
FIPS-CC mode: disable
Current HA mode: standalone
Distribution: International
Branch point: 511
Release Version Information: MR3 Patch 4
System time: Wed Jan 18 11:24:34 2012

diagnose hardware sysinfo vm full
    The following output is displayed:
UUID: 564db33a29519f6b1025bf8539a41e92
valid: 1
status: 1
code: 200 (If the license is a duplicate, code 401 will be displayed)
warn: 0
copy: 0
received: 45438
warning: 0
recv: 201201201918
dup:

```

Licensing timeout

In closed environments without Internet access, it is mandatory to perform offline licensing of the virtual FortiGate using a FortiManager as a license server. If the FortiGate-VM cannot perform license validation within the license timeout period, which is 30 days, the FortiGate will discard all packets and effectively ceasing operation as a firewall.

The status of the license will go through some status changes before it times out.

Status	Description
Valid	The FortiGate can connect and validate against a FortiManager or FDS
Warning	The FortiGate cannot connect and validate against a FortiManager or FDS. A check is made against how many days the Warning status has been continuous. If the number is less the 30 days the status does not change.
Invalid	The FortiGate cannot connect and validate against a FortiManager or FDS. A check is made against how many days the Warning status has been continuous. If the number is 30 days or more, the status changes to Invalid. The firewall ceases to function properly.



There is only a single log entry after the virtual FortiGate cannot access the license server for the license expiration period. This means that when you go searching the logs for a reason for the FortiGate being offline there will not be a long list of error logs that draw attention to the issue. There will only be the one entry.

Configure your FortiGate VM

Once the FortiGate VM license has been validated you can begin to configure your device. You can use the **Wizard** located in the top toolbar for basic configuration including enabling central management, setting the admin password, setting the time zone, and port configuration.

For more information on configuring your FortiGate VM see the FortiOS Handbook at <http://docs.fortinet.com>.

Chapter 30 - VoIP Solutions: SIP

This FortiOS Handbook chapter contains the following sections:

[Inside FortiOS: VoIP Protection](#) introduces FortiOS VoIP Protection

[Common SIP VoIP configurations](#) describes some common SIP configurations.

[SIP messages and media protocols](#) describes SIP messages and some common SIP media protocols.

[The SIP session helper](#) describes how the SIP session helper works and how to configure SIP support using the SIP session helper.

[The SIP ALG](#) describes how the SIP Application Layer Gateway (ALG) works and how to configure SIP support using the SIP ALG.

[Conflicts between the SIP ALG and the session helper](#) describes how to sort out conflicts between the SIP session helper and the ALG.

[Stateful SIP tracking, call termination, and session inactivity timeout](#) describes how the SIP ALG performs SIP stateful tracking, call termination and session activity timeouts.

[SIP and RTP/RTCP](#) describes how SIP relates to RTP and RTCP.

[How the SIP ALG creates RTP pinholes](#) describes how the SIP ALG creates pinholes.

[Configuration example: SIP in transparent mode](#) describes how to configure a FortiGate in transparent mode to support SIP.

[RTP enable/disable \(RTP bypass\)](#) describes RTP bypass.

[Opening and closing SIP register, contact, via and record-route pinholes](#) describes how FortiOS opens and closes these pinholes.

[Accepting SIP register responses](#) describes how to enable accepting SIP register responses.

[How the SIP ALG performs NAT](#) describes how the SIP ALG performs NAT.

[Enhancing SIP pinhole security](#) describes how to open smaller pinholes.

[Hosted NAT traversal](#) describes SIP hosted NAT traversal and how to configure it.

[SIP over IPv6](#) describes how to configure SIP over IPv6.

[Deep SIP message inspection](#) describes how deep SIP message inspection works.

[Blocking SIP request messages](#) describes how to block SIP request messages to prevent some common SIP attacks.

[SIP rate limiting](#) includes more options for preventing SIP attacks.

[SIP logging](#) describes how to enable SIP logging.

[Inspecting SIP over SSL/TLS \(secure SIP\)](#) describes how to inspect encrypted SIP traffic.

[SIP and HA-session failover and geographic redundancy](#) describes how to use FGCP HA to support SIP geographic redundancy.

[SIP and IPS](#) describes how to turn on IPS for SIP sessions.

[SIP debugging](#) describes some tools for debugging your SIP configuration.

What's new in FortiOS 6.0.1

VoIP features appear on the GUI when the FortiGate is operating in Flow mode, see [Enabling VoIP support from the GUI on page 3124](#).

What's new in FortiOS 6.0

By default, FortiOS 6.0 disables the SIP session helper, see [SIP session helper configuration overview on page 3116](#).

Inside FortiOS: Voice over IP (VoIP) protection

The FortiOS SIP Application Layer Gateway (ALG) allows SIP calls to pass through a FortiGate by opening SIP and RTP pinholes and performing source and destination IP address and port translation for SIP and RTP packets.

There are a large number of SIP-related Internet Engineering Task Force (IETF) documents (Request for Comments) that define behavior of SIP and related applications. FortiOS completely support [RFC 3261](#) for SIP, [RFC 4566](#) for SDP and [RFC 3262](#) for Provisional Response Acknowledgment (PRACK). FortiOS also supports other SIP and SIP-related RFCs and performs Deep SIP message inspection for SIP statements defined in other SIP RFCs.

Advanced voice over IP protection


The FortiOS SIP Application Level Gateway (ALG) protects Voice over IP (SIP and SDP) services in Unified Communication and NGN/IMS networks with the following advanced VoIP defense mechanisms.

Deep SIP message inspection (also called deep SIP header inspection)

Verifies SIP and SDP header syntax and protects SIP servers from potential SIP Fuzzing attacks. When a violation is detected, FortiOS can impose counter measures and can also send automatic SIP response messages to offload processing from the SIP server.

SIP message rate limiting

Allows rate limiting of SIP messages per SIP request method. This prevents a SIP server from overload or from DoS attacks using particular SIP methods. For example, FortiOS can protect SIP servers from a flood of SIP REGISTER or INVITE messages, which can be caused by a DoS attack or a flash crowd.



RTP and RTCP pinholing

RTP pinholing only forwards RTP/RTCP packets that conform to the particular session description of the associated SIP dialog. If a SIP dialog is finished, FortiOS automatically closes the pinhole. RTP/RTCP pinholing is supported by FortiASIC acceleration and achieves high packet throughput at low jitter and delay.

Stateful SIP dialog tracking

FortiOS tracks SIP message sequences and prevents unwanted SIP messages that are not related to a particular SIP dialog. For instance, FortiOS can detect malicious SIP BYE messages that do not conform with the associated context of the SIP dialog.

Inspecting SIP over SSL/TLS (secure SIP)

Some SIP phones and SIP servers use SSL or TLS to encrypt SIP signalling traffic. To allow SIP over SSL/TLS calls to pass through the FortiGate unit, the encrypted signalling traffic has to be unencrypted and inspected. FortiOS intercepts and unencrypts and inspects the SIP packets. Allowed packets are then re-encrypted and forwarded to their destination.

Inspecting SIP on multiple ports

FortiOS can detect and inspect SIP and SDP UDP and TCP sessions and SIP SSL sessions and you can configure the ports that the SIP ALG monitors for these sessions. In addition you can configure two different ports for SIP UDP sessions and two different ports for SIP TCP sessions. The port configuration can be changed without affecting other parts of the SIP configuration.

Carrier grade protection

To protect VoIP infrastructure in carrier networks, FortiOS complies with typical carrier requirements for availability and robustness.

High availability

FortiOS supports a hot failover configuration with an active and a standby FortiGate device. FortiOS dynamically updates the context on the standby unit with SIP and RTP related data. This enables the standby unit to takeover stable voice calls in case of a planned or unplanned outage or failover of the active unit.

Geographical redundancy of SIP servers

In FortiOS SIP server cluster configurations the active and standby units can be deployed in different geographical locations. This configuration prevents a total outage of a SIP server infrastructure if one location goes offline. FortiOS supports the detection of SIP server outages (loss of heartbeats) and a redirect of SIP messages to the redundant SIP server location.

Logging and Reporting

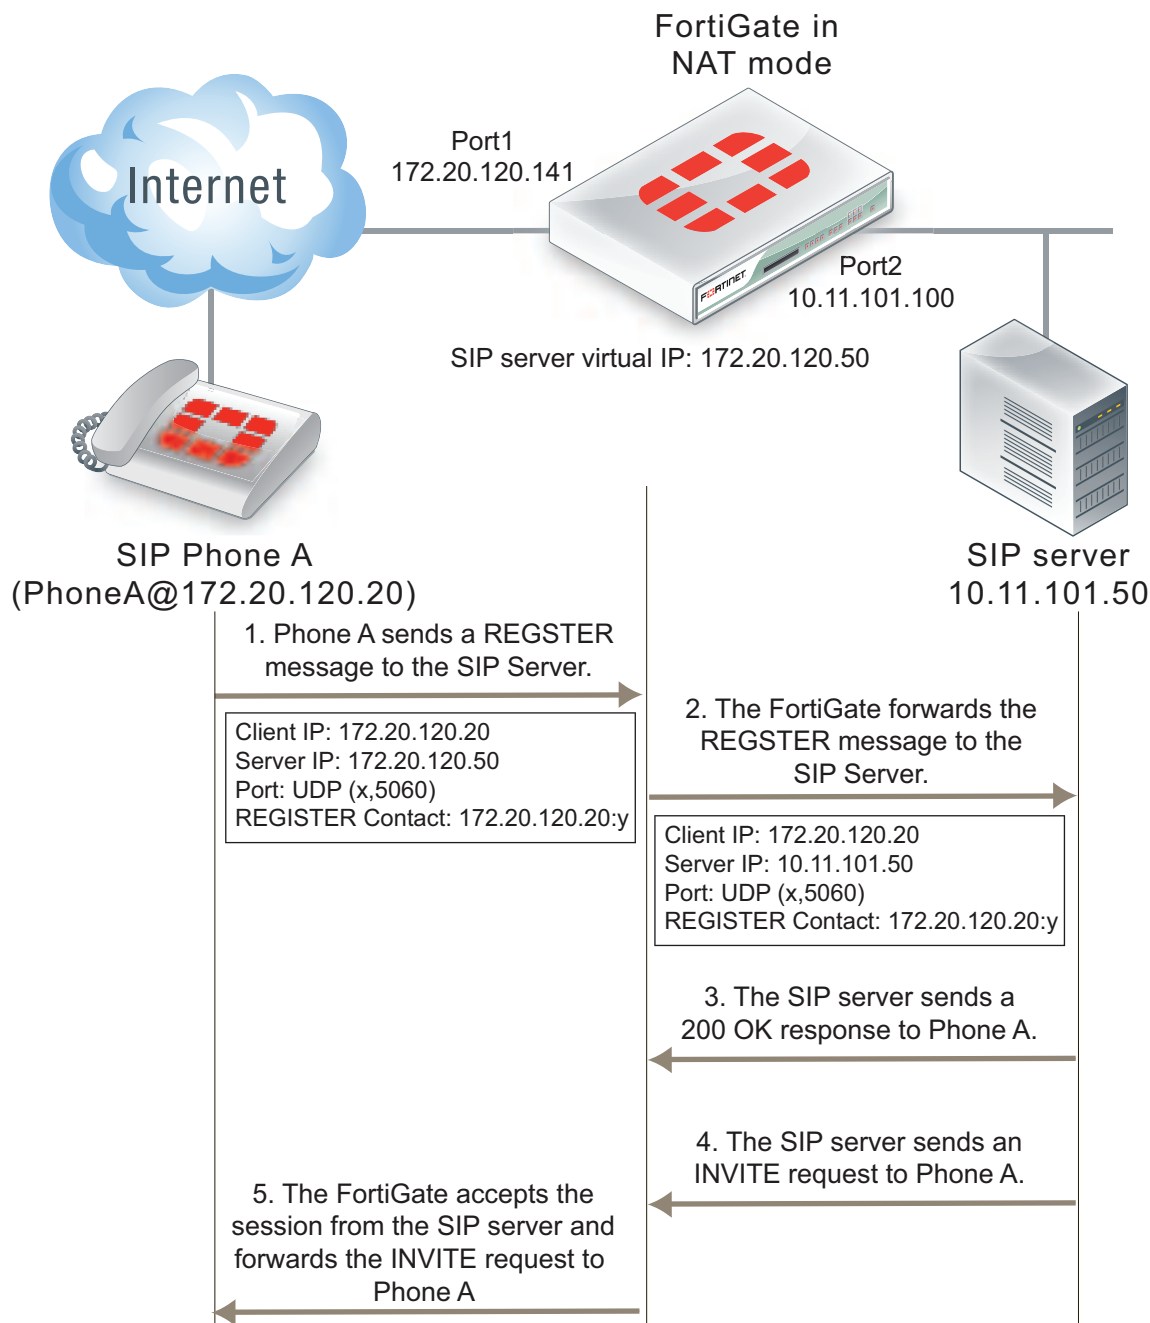
FortiOS can log call related information internally or to an external SYSLOG or FortiAnalyzer unit. This includes event logs that show particular SIP-related attacks or syntax violations with SIP messages or logs that summarize call statistics.

NAT/NAPT

FortiOS performs configurable network address translation for IP addresses in the SIP and SDP header. The SIP ALG follows the configured NAT addresses in firewall virtual IPs and changes SIP header IP addresses accordingly. RTP NAT is controlled by SIP/SDP and the firewall policy. This allows translating an unlimited number of IP addresses without adding specific RTP policies.

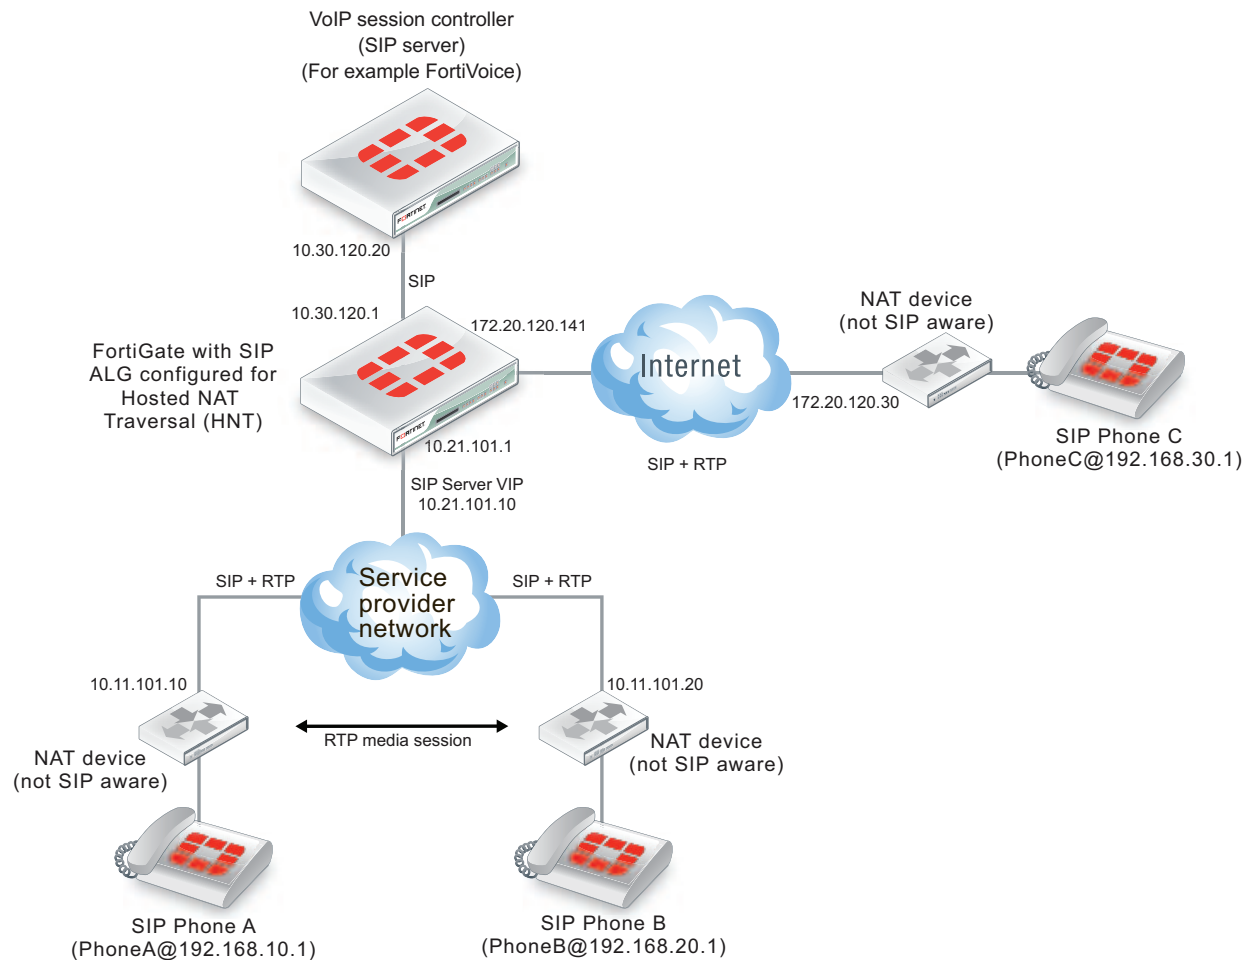
Header manipulation

FortiOS SIP and SDP header manipulation supports SIP Network Address Translation (NAT) through FortiGate units configured as NAT firewalls.



Hosted NAT traversal (HNT)

In many service provider networks, CPE firewall devices provide NAT without application awareness. This causes issues for SIP/SDP and RTP traffic, since UAC IP address information references to the internal network behind the far end firewall. VoIP calls cannot be connected successfully. FortiOS mitigates far end NAT issues (called Hosted NAT traversal) by probing the first RTP packet from the UAC and learning the far end NA(P)T binding. FortiOS then updates the internal NAT binding for RTP accordingly.

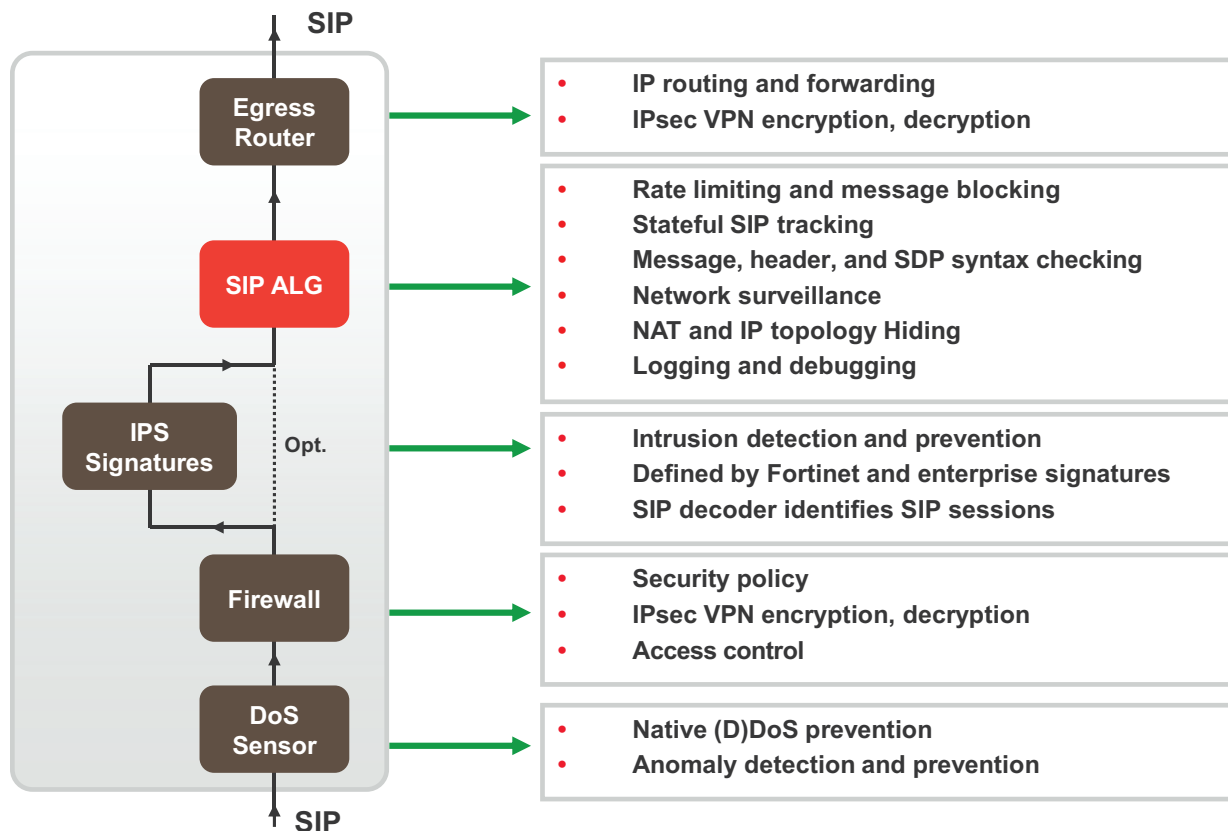


IP address conservation for NAT

In case of SIP and RTP NAT IP the original address information can get lost after translating to the provisioned IP addresses. This IP address information is sometimes required for detailed billing records or debugging purposes. FortiOS can maintain the original IP address information in a translated SIP header by adding it to the SIP/SDP info line (i=) or by adding it to the original attribute (o=). Either option can be selected depending on the SIP billing environment.

SIP ALG activation

The FortiOS SIP ALG is applied to SIP traffic accepted by a firewall policy that includes a VoIP profile. The VoIP profile controls how the SIP ALG processes SIP sessions. FortiOS also includes a high-performance SIP session helper that provides limited SIP functionality. In most cases the SIP ALG should be used because the SIP ALG supports the complete range of FortiOS SIP features.



SIP over IPv6

FortiOS, operating in NAT/Route and in transparent mode supports SIP over IPv6. The SIP ALG can process SIP messages that use IPv6 addresses in the headers, bodies, and in the transport stack. The SIP ALG cannot modify the IPv6 addresses in the SIP headers so FortiGate units cannot perform SIP or RTP NAT over IPv6 and also cannot translate between IPv6 and IPv4 addresses.

Platform support and hardware acceleration

FortiOS supports VoIP protection with the SIP ALG on all FortiGate hardware platforms. Whenever a FortiGate unit provides FortiASIC or SPM HW acceleration, the SIP ALG will use this option to fast-path RTP/RTCP traffic.

As well, since the SIP ALG is proxy-based, SIP control packets are not offloaded to NP4 or NP6 processors. But actual voice or other media traffic can be offloaded to NP4 or NP6 processors after the SIP session is established. Many FortiGate units also support low latency hardware acceleration configurations that also enhance SIP voice transmission.

FortiGate hardware acceleration provides a high throughput solution at very low jitter and delay. FortiOS provides efficient and highly scalable protection for VoIP in emerging Enterprise and Carrier network. This complements Fortinet's NGFW and UTM offerings. VoIP protection can be easily added to any firewall policy just by adding a VoIP profile.

VoIP protection is supported in FortiAnalyzer and FortiManager. Centralized logging and management are essential for carrier and MSSP service provider and are influencing business case calculations.

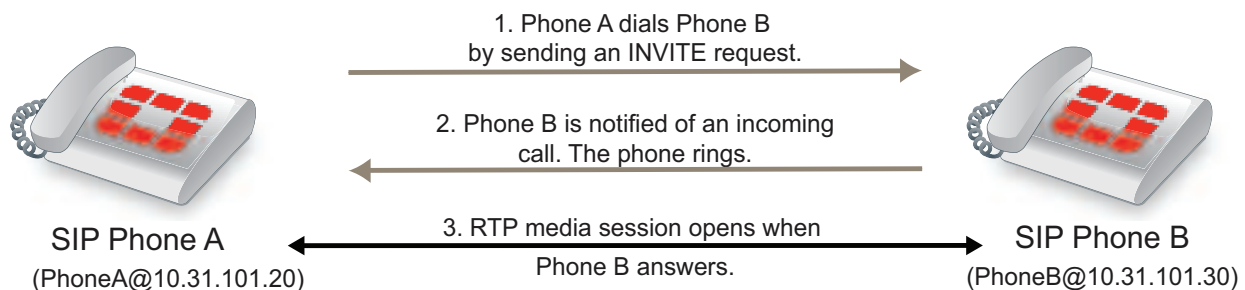
Common SIP VoIP configurations

This section describes some common SIP VoIP configurations and simplified SIP dialogs for these configurations. This section also shows some examples of how adding a FortiGate affects SIP processing.

Peer to peer configuration

In the peer to peer configuration shown below, two SIP phones (in the example, FortiFones) communicate directly with each other. The phones send SIP request and response messages back and forth between each other to establish the SIP session.

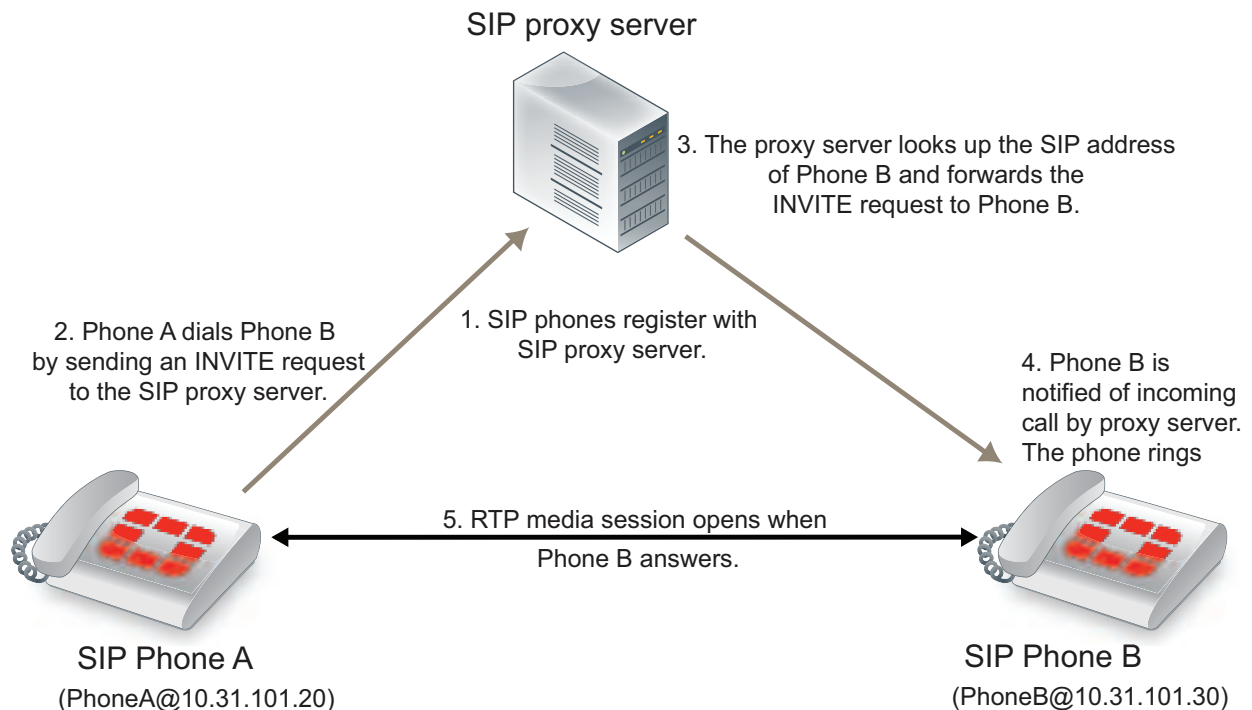
SIP peer to peer configuration



Peer to peer configurations are not very common because they require the SIP phones to keep track of the names and addresses of all of the other SIP phones that they can communicate with. In most cases a SIP proxy or re-direct server maintains addresses of a large number of SIP phones and a SIP phone starts a call by contacting the SIP proxy server.

SIP proxy server configuration

A SIP proxy server acts as an intermediary between SIP phones and between SIP phones (for example, two FortiFones) and other SIP servers. As shown below, SIP phones send request and response messages to the SIP proxy server. The proxy server forwards the messages to other clients or to other SIP proxy servers. Proxy servers can hide SIP phones by proxying the signaling messages. To the other users on the VoIP network, the signaling invitations look as if they come from the SIP proxy server.

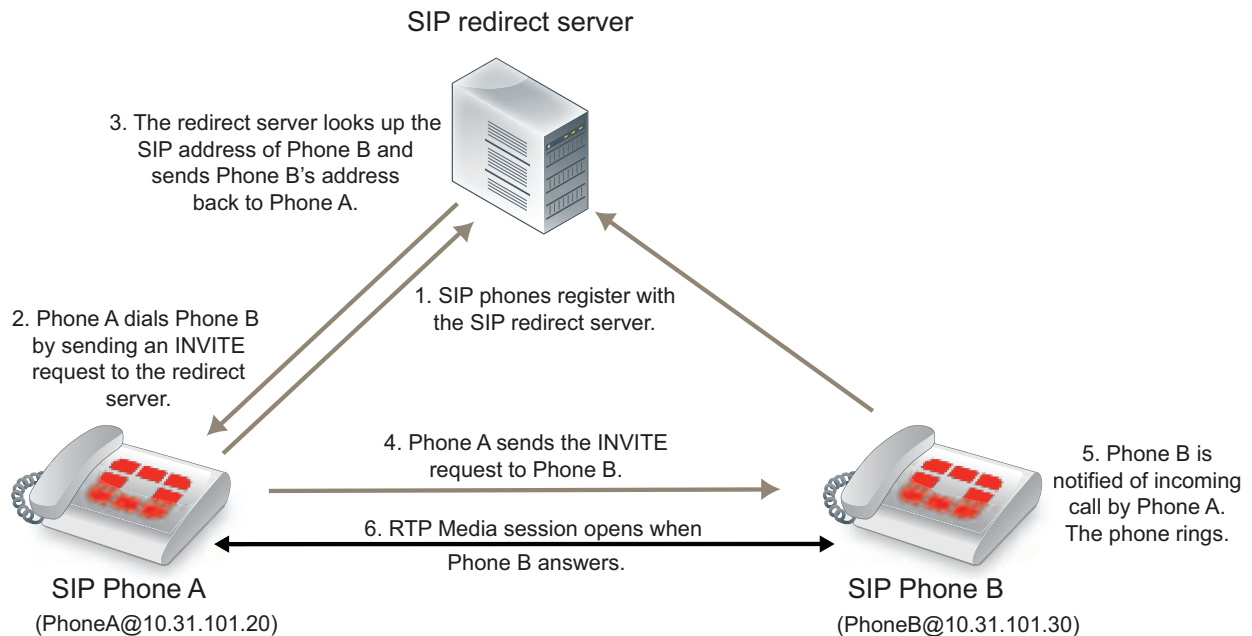
SIP in proxy mode

A common SIP configuration would include multiple networks of SIP phones. Each of the networks would have its own SIP server. Each SIP server would proxy the communication between phones on its own network and between phones in different networks.

SIP redirect server configuration

A SIP redirect server accepts SIP requests, maps the addresses in the request into zero or more new addresses and returns those addresses to the client. The redirect server does not initiate SIP requests or accept calls. As shown below, SIP clients send INVITE requests to the redirect server, which then looks up the destination address. The redirect server returns the destination address to the client. The client uses this address to send the INVITE request directly to the destination SIP client.


SIP in redirect model



SIP registrar configuration

A SIP registrar accepts SIP REGISTER requests from SIP phones for the purpose of updating a location database with this contact information. This database can then become a SIP location service that can be used by SIP proxy servers and redirect servers to locate SIP clients. As shown below, SIP clients send REGISTER requests to the SIP registrar.

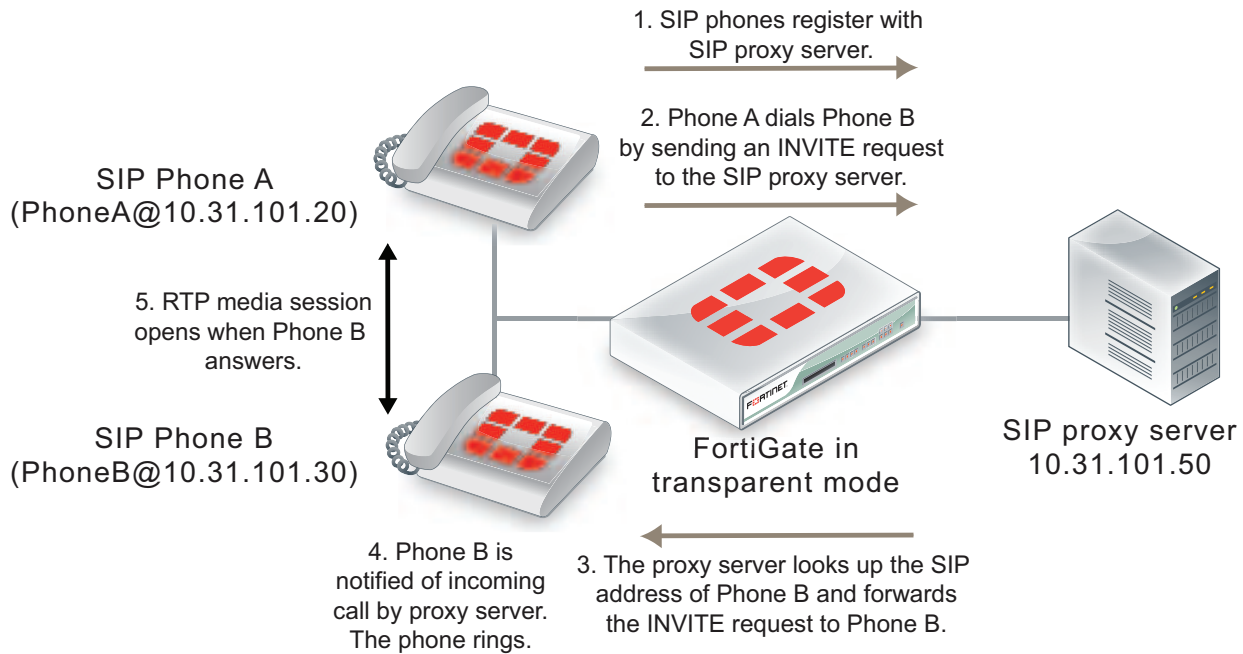
SIP registrar and proxy servers



SIP with a FortiGate

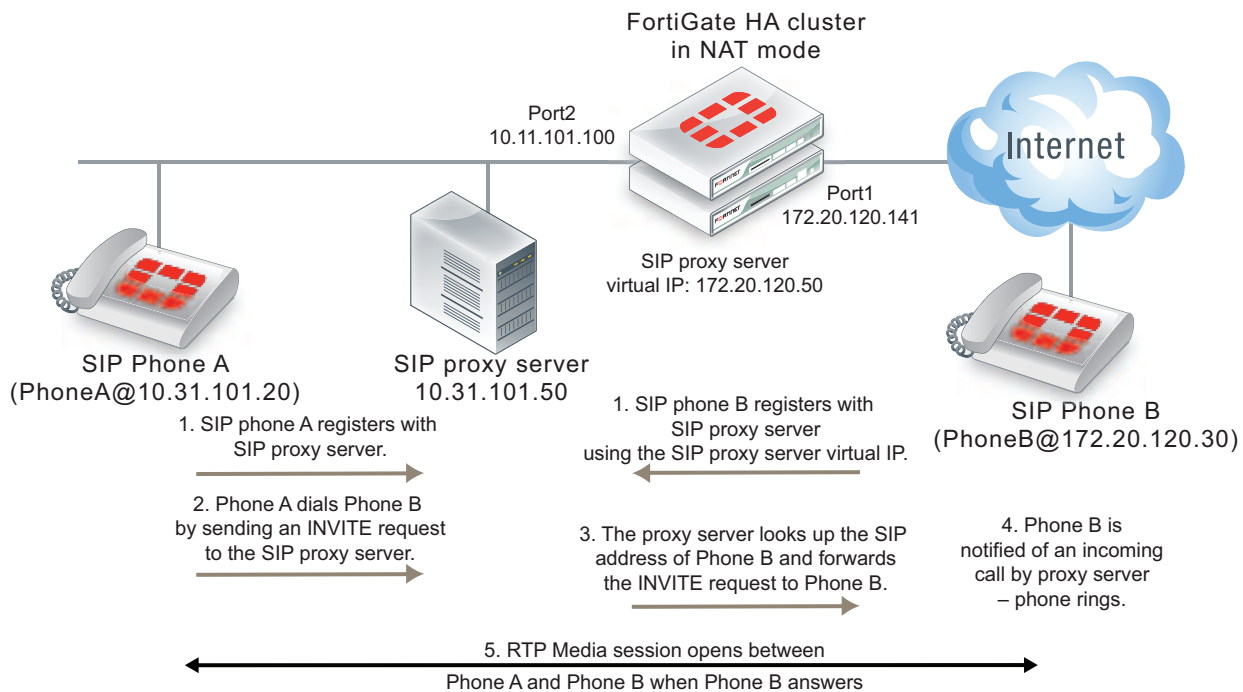
Depending on your security requirements and network configuration FortiGates may be in many different places in a SIP configuration. This section shows a few examples.

The diagram below shows a FortiGate installed between a SIP proxy server and SIP phones on the same network. The FortiGate is operating in transparent mode so both the proxy server and the phones are on the same subnet. In this configuration, called SIP inspection without address translation, the FortiGate could be protecting the SIP proxy server on the private network by implementing SIP security features for SIP sessions between the SIP phones and the SIP proxy server.

SIP network with FortiGate in transparent mode

The phones and server use the same SIP dialogs as they would if the FortiGate was not present. However, the FortiGate can be configured to control which devices on the network can connect to the SIP proxy server and can also protect the SIP proxy server from SIP vulnerabilities.

The following diagram shows a FortiGate operating in NAT/Route mode and installed between a private network and the Internet. Some SIP phones and the SIP proxy server are connected to the private network and some SIP phones are connected to the Internet. The SIP phones on the Internet can connect to the SIP proxy server through the FortiGate and communication between SIP phones on the private network and SIP phones on the Internet must pass through the FortiGate.

SIP network with FortiGate in NAT/Route mode

The phones and server use the same SIP dialog as they would if the FortiGate was not present. However, the FortiGate can be configured to control which devices on the network can connect to the SIP proxy server and can also protect the SIP proxy server from SIP vulnerabilities. In addition, the FortiGate has a firewall virtual IP that forwards packets sent to the SIP proxy server Internet IP address (172.20.120.50) to the SIP proxy server internal network IP address (10.31.101.30).

Since the FortiGate is operating in NAT/Route mode it must translate packet source and destination IP addresses (and optionally ports) as the sessions pass through the FortiGate. Also, the FortiGate must translate the addresses contained in the SIP headers and SDP body of the SIP messages. As well the FortiGate must open SIP and RTP pinholes through the FortiGate. SIP pinholes allow SIP signaling sessions to pass through the FortiGate between phones and between phones and SIP servers. RTP pinholes allow direct RTP communication between the SIP phones once the SIP dialog has established the SIP call. Pinholes are opened automatically by the FortiGate. Administrators do not add security policies for pinholes or for RTP sessions. All that is required is a security policy that accepts SIP traffic.

Opening an RTP pinhole means opening a port on a FortiGate interface to allow RTP traffic to use that port to pass through the FortiGate between the SIP phones on the Internet and SIP phones on the internal network. A pinhole only accepts packets from one RTP session. Since a SIP call involves at least two media streams (one from Phone A to Phone B and one from Phone B to Phone A) the FortiGate opens two RTP pinholes. Phone A sends RTP packets through a pinhole in port2 and Phone B sends RTP packets through a pinhole in port1. The FortiGate opens the pinholes when required by the SIP dialog and closes the pinholes when the SIP call is completed. The FortiGate opens new pinholes for each SIP call.

Each RTP pinhole actually includes two port numbers. The RTP port number as defined in the SIP message and an RTCP port number, which is the RTP port number plus 1. For example, if the SIP call used RTP port 3346 the FortiGate would create a pinhole for ports 3346 and 3347.

SIP messages and media protocols

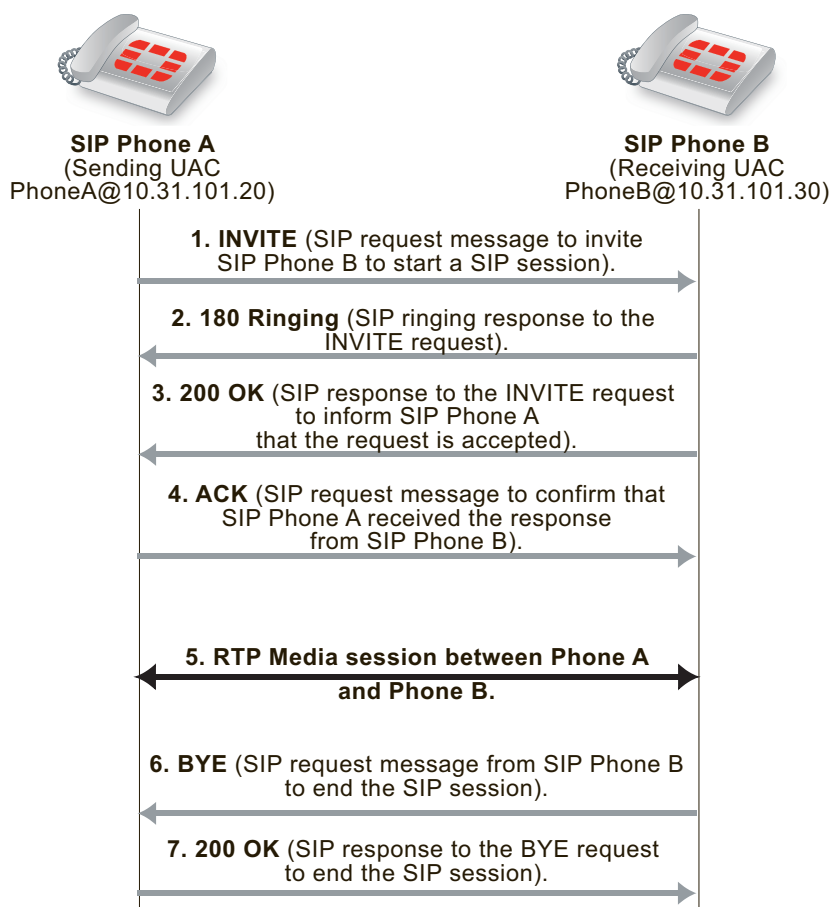
This section provides an overview of SIP messages and how they communicate information about SIP sessions and how SDP, RTP, and RTCP fits in with SIP communications.

SIP uses clear text messages to start, maintain, and end media sessions between SIP user agent clients (UACs) and user agent servers (UASs). These messages form a SIP dialog. A typical SIP dialog begins with an INVITE request message sent from a UAC to another UAC or to a UAS. The first INVITE request message attempts to start a SIP call and includes information about the sending UAC and the receiving UAC as well as information about the communication session.

If only two UACs are involved as shown below, the receiving UAC (Phone B) responds with a 180 Ringing and then a 200 OK SIP response message that informs Phone A that Phone B received and accepted the request. Phone A then sends an ACK message to notify Phone B that the SIP response was received. Phone A and Phone B can then participate in the RTP media session set up by the SIP messages.

When the phone call is complete, one of the UACs (in the example Phone B) hangs up sending a BYE request message to Phone A. Phone A then sends a 200 OK response to Phone B acknowledging that the session has ended.

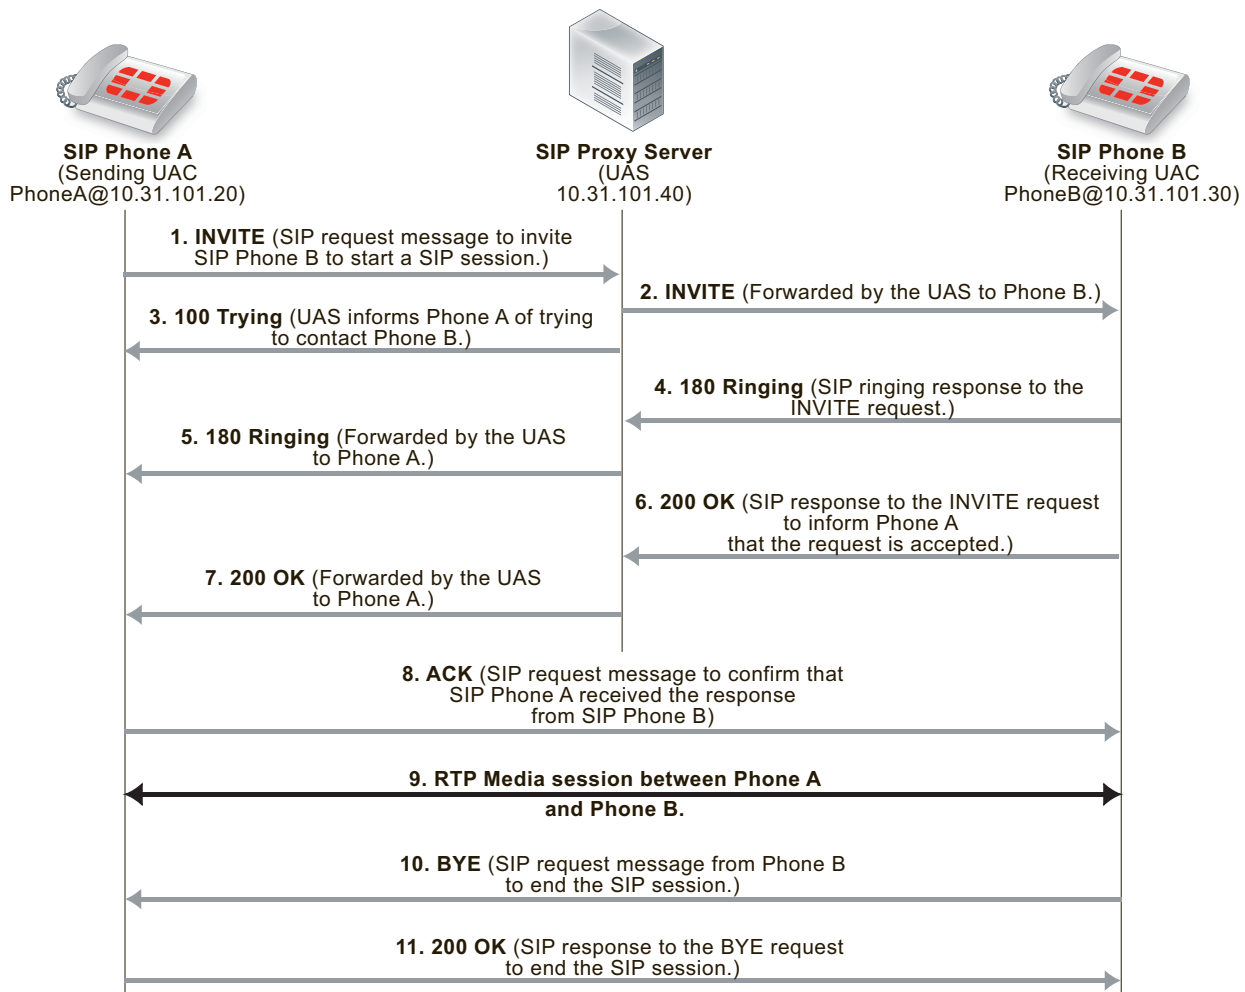
Basic SIP dialog between two UACs



If a UAS in the form of a SIP proxy server is involved, similar messages are sent and received, but the proxy server participates as an intermediary in the initial call setup. In the example below the SIP proxy server receives the INVITE request from Phone A and forwards it to Phone B. The proxy server then sends a 100 Trying response to Phone A. Phone B receives the INVITE request and responds with a 180 Ringing and then a 200 OK SIP response message. These messages are received by the proxy server and forwarded to Phone A to notify Phone A that Phone B received and accepted the request. Phone A then sends an ACK message to notify Phone B that the SIP response was received. This response is received by the proxy server and forwarded to Phone B. Phone A and Phone B can then participate in the media session independently of the proxy server.

When the phone call is complete Phone B hangs up sending a BYE request message to Phone A. Phone A then sends a 200 OK response to Phone B acknowledging that the session has ended.

Basic SIP dialog between UACs with a SIP proxy server UAS



The SIP messages include SIP headers that contain names and addresses of Phone A, Phone B and the proxy server. This addressing information is used by the UACs and the proxy server during the call set up.

The SIP message body includes Session Description Protocol (SDP) statements that Phone A and Phone B use to establish the media session. The SDP statements specify the type of media stream to use for the session (for example, audio for SIP phone calls) and the protocol to use for the media stream (usually the Real Time Protocol (RTP) media streaming protocol).

Phone A includes the media session settings that it would like to use for the session in the INVITE message. Phone B includes its response to these media settings in the 200 OK response. Phone A's ACK response confirms the settings that Phone A and Phone B then use for the media session.

Hardware accelerated RTP processing

FortiGates can offload RTP packet processing to network processor (NP) interfaces. This acceleration greatly enhances the overall throughput and resulting in near speed RTP performance.

SIP request messages

SIP sessions always start with a SIP request message (also just called a SIP request). SIP request messages also establish, maintain, and terminate SIP communication sessions. The following table lists some common SIP request message types.

Common SIP request message types

Message Type	Description
INVITE	A client sends an INVITE request to invite another client to participate in a multimedia session. The INVITE request body usually contains the description of the session.
ACK	The originator of an INVITE message sends an ACK request to confirm that the final response to an INVITE request was received. If the INVITE request did not contain the session description, it must be included in the ACK request.
PRACK	In some cases, SIP uses provisional response messages to report on the progress of the response to a SIP request message. The provisional response messages are sent before the final SIP response message. Similar to an ACK request message, a PRACK request message is sent to acknowledge that a provisional response message has been received.
OPTIONS	The UA uses OPTIONS messages to get information about the capabilities of a SIP proxy. The SIP proxy server replies with a description of the SIP methods, session description protocols, and message encoding that are supported.
BYE	A client sends a BYE request to end a session. A BYE request from either end of the SIP session terminates the session.
CANCEL	A client sends a CANCEL request to cancel a previous INVITE request. A CANCEL request has no effect if the SIP server processing the INVITE sends a final response to the INVITE before receiving the CANCEL.
REGISTER	A client sends a REGISTER request to a SIP registrar server with information about the current location (IP address and so on) of the client. A SIP registrar server saves the information it receives in REGISTER requests and makes this information available to any SIP client or server attempting to locate the client.

Message Type	Description
Info	For distributing mid-session signaling information along the signaling path for a SIP call. I
Subscribe	For requesting the current state and state updates of a remote node.
Notify	Informs clients and servers of changes in state in the SIP network.
Refer	Refers the recipient (identified by the Request-URI) to a third party according to the contact information in the request.
Update	Opens a pinhole for new or updated SDP information.
Response codes (1xx, 202, 2xx, 3xx, 4xx, 5xx, 6xx)	Indicates the status of a transaction. For example: 200 OK, 202 Accepted, or 400 Bad Request.

SIP response messages

SIP response messages (often just called SIP responses) provide status information in response to SIP request messages. All SIP response messages include a response code and a reason phrase. There are five SIP response message classes. They are described below.

There are also two types of SIP response messages, provisional and final. Final response messages convey the result of the request processing, and are sent reliably. Provisional responses provide information on the progress of the request processing, but may not be sent reliably. Provisional response messages start with 1xx and are also called informational response messages.

Informational (or provisional)

Informational or provisional responses indicate that a request message was received and imply that the endpoint is going to process the request. Information messages may not be sent reliably and may not require an acknowledgment.

If the SIP implementation uses Provisional Response Acknowledgment (PRACK) ([RFC 3262](#)) then informational or provisional messages are sent reliably and require a PRACK message to acknowledge that they have been received.

Informational responses can contain the following reason codes and reason phrases:

```
100 Trying
180 Ringing
181 Call is being forwarded
182 Queued
183 Session progress
```

Success

Success responses indicate that a request message was received, understood, and accepted. Success responses can contain the following reason codes and reason phrases:

```
200 OK
```

202 Accepted

Redirection

Redirection responses indicate that more information is required for the endpoint to respond to a request message. Redirection responses can contain the following reason codes and reason phrases:

300 Multiple choices
301 Moved permanently
302 Moved temporarily
305 Use proxy
380 Alternative service

Client error

Client error responses indicate that a request message was received by a server that contains syntax that the server cannot understand (i.e. contains a syntax error) or cannot comply with. Client error responses include the following reason codes and reason phrases:

400 Bad request	401 Unauthorized
402 Payment required	403 Forbidden
404 Not found	405 Method not allowed
406 Not acceptable	407 Proxy authentication required
408 Request time-out	409 Conflict
410 Gone	411 Length required
413 Request entity too large	414 Request-URL too large
415 Unsupported media type	420 Bad extension
480 Temporarily not available	
481 Call leg/transaction does not exist	
482 Loop detected	
484 Address incomplete	483 Too many hops
486 Busy here	485 Ambiguous
488 Not acceptable here	487 Request canceled

Server error

Server error responses indicate that a server was unable to respond to a valid request message. Server error responses include the following reason codes and reason phrases:

500 Server internal error
501 Not implemented
502 Bad gateway
502 Service unavailable
504 Gateway time-out
505 SIP version not supported

Global failure

Global failure responses indicate that there are no servers available that can respond to a request message. Global failure responses include the following reason codes and reason phrases:

600 Busy everywhere
603 Decline
604 Does not exist anywhere
606 Not acceptable

SIP message start line

The first line in a SIP message is called the start line. The start line in a request message is called the request-line and the start line in a response message is called the status-line.

Request-line	<p>The first line of a SIP request message. The request-line includes the SIP message type, the SIP protocol version, and a Request URI that indicates the user or service to which this request is being addressed. The following example request-line specifies the INVITE message type, the address of the sender of the message (inviter@example.com), and the SIP version:</p> <pre>INVITE sip:inviter@example.com SIP/2.0</pre>
Status-line	<p>The first line of a SIP response message. The status-line includes the SIP protocol version, the response code, and the reason phrase. The example status-line includes the SIP version, the response code (200) and the reason phrase (OK).</p> <pre>SIP/2.0 200 OK</pre>

SIP headers

Following the start line, SIP messages contain SIP headers (also called SIP fields) that convey message attributes and to modify message meaning. SIP headers are similar to HTTP header fields and always have the following format:

```
<header_name>:<value>
```

SIP messages can include the SIP headers listed in the following table:

SIP headers

SIP Header	Description
Allow	<p>Lists the set of SIP methods supported by the UA generating the message. All methods, including ACK and CANCEL, understood by the UA MUST be included in the list of methods in the Allow header field, when present. For example:</p> <pre>Allow: INVITE, ACK, OPTIONS, CANCEL, BYE</pre>
Call-ID	<p>A globally unique identifier for the call, generated by the combination of a random string and the sender's host name or IP address. The combination of the To, From, and Call-ID headers completely defines a peer-to-peer SIP relationship between the sender and the receiver. This relationship is called a SIP dialog.</p> <pre>Call-ID: ddeg45e793@10.31.101.30</pre>

SIP Header	Description
Contact	<p>Included in SIP request messages, the Contact header contains the SIP URI of the sender of the SIP request message. The receiver uses this URI to contact the sender. For example:</p> <pre>Contact: Sender <sip:sender@10.31.100.20>t</pre>
Content-Length	<p>The number of bytes in the message body (in bytes).</p> <pre>Content-Length: 126</pre>
Content-Type	<p>In addition to SIP headers, SIP messages include a message body that contains information about the content or communication being managed by the SIP session. The Content-Type header specifies what the content of the SIP message is. For example, if you are using SIP with SDP, the content of the SIP message is SDP code.</p> <pre>Content-Type: application/sdp</pre>
CSeq	<p>The command sequence header contains a sequence integer that is increased for each new SIP request message (but is not incremented in the response message). This header also includes the request name found in the request message request-line. For example:</p> <pre>CSeq: 1 INVITE</pre>
Expires	<p>Gives the relative time after which the message (or content) expires. The actual time and how the header is used depends on the SIP method. For example:</p> <pre>Expires: 5</pre>
From	<p>Identifies the sender of the message. Responses to a message are sent to the address of the sender. The following example includes the sender's name (<i>Sender</i>) and the sender's SIP address (<i>sender@10.31.101.20</i>):</p> <pre>From: Sender <sip:sender@10.31.101.20></pre>
Max-forwards	<p>An integer in the range 0-255 that limits the number of proxies or gateways that can forward the request message to the next downstream server. Also called the number of hops, this value is decreased every time the message is forwarded. This can also be useful when the client is attempting to trace a request chain that appears to be failing or looping in mid-chain.</p> <pre>For example: Max-Forwards: 30</pre>
P-Asserted-Identity	<p>The P-Asserted-Identity header is used among trusted SIP entities to carry the identity of the user sending a SIP message as it was verified by authentication. See RFC 3325. The header contains a SIP URI and an optional display-name, for example:</p> <pre>P-Asserted-Identity: "Example Person" <sip:10.31.101.50></pre>

SIP Header	Description
RAck	<p>Sent in a PRACK request to support reliability of information or provisional response messages. It contains two numbers and a method tag. For example:</p> <pre>RAck: 776656 1 INVITE</pre>
Record-Route	<p>Inserted into request messages by a SIP proxy to force future requests to be routed through the proxy. In the following example, the host at IP address 10.31.101.50 is a SIP proxy. The <code>lr</code> parameter indicates the URI of a SIP proxy in Record-Route headers.</p> <pre>Record-Route: <sip:10.31.101.50;lr></pre>
Route	<p>Forces routing for a request message through one or more SIP proxies. The following example includes two SIP proxies:</p> <pre>Route: <sip:172.20.120.10;lr>, <sip:10.31.101.50;lr></pre>
RSeq	<p>The RSeq header is used in information or provisional response messages to support reliability of informational response messages. The header contains a single numeric value. For example:</p> <pre>RSeq: 33456</pre>
To	<p>Identifies the receiver of the message. The address in this field is used to send the message to the receiver. The following example includes the receiver's name (<code>Receiver</code>) and the receiver's SIP address (<code>receiver@10.31.101.30</code>):</p> <pre>To: Receiver <sip:receiver@10.31.101.30></pre>
Via	<p>Indicates the SIP version and protocol to be used for the SIP session and the address to which to send the response to the message that contains the Via field. The following example Via field indicates to use SIP version 2, UDP for media communications, and to send the response to 10.31.101.20 using port 5060.</p> <pre>Via: SIP/2.0/UDP 10.31.101.20:5060</pre>

The SIP message body and SDP session profiles

The SIP message body describes the session to be initiated. For example, in a SIP phone call the body usually includes audio codec types, sampling rates, server IP addresses and so on. For other types of SIP session the body could contain text or binary data of any type which relates in some way to the session. The message body is included in request and response messages.

Two possible SIP message body types:

- Session Description Protocol (SDP), most commonly used for SIP VoIP.
- Multipurpose Internet Mail Extensions (MIME)

SDP is most often used for VoIP and FortiGates support SDP content in SIP message bodies. SDP is a text-based protocol used by SIP to control media sessions. SDP does not deliver media but provides a session profile

that contains media details, transport addresses, parameter negotiation, and other session description metadata for the participants in a media session. The participants use the information in the session profile to negotiate how to communicate and to manage the media session. SDP is described by [RFC 4566](#).

An SDP session profile always contains session information and may contain media information. Session information appears at the start of the session profile and media information (using the `m=` attribute) follows.

SDP session profiles can include the attributes listed in the following table.

SDP session profile attributes

Attribute	Description
a=	Attributes to extend SDP in the form <code>a=<attribute></code> or <code>a=<attribute>:<value></code> .
b=	Contains information about the bandwidth required for the session or media in the form <code>b=<bandwidth_type>:<bandwidth></code> .
c=	Connection data about the session including the network type (usually IN for Internet), address type (IPv4 or IPv6), the connection source address, and other optional information. For example: <code>c=IN IPv4 10.31.101.20</code>
i=	A text string that contains information about the session. For example: <code>i=A audio presentation about SIP</code>
k=	Can be used to convey encryption keys over a secure and trusted channel. For example: <code>k=clear:444gdduudjffdee</code>

Attribute	Description
m=	<p>Media information, consisting of one or more lines all starting with m= and containing details about the media including the media type, the destination port or ports used by the media, the protocol used by the media, and a media format description.</p> <pre>m=audio 49170 RTP 0 3 m-video 3345/2 udp 34 m-video 2910/2 RTP/AVP 3 56</pre> <p>Multiple media lines are needed if SIP is managing multiple types of media in one session (for example, separate audio and video streams).</p> <p>Multiple ports for a media stream are indicated using a slash. <code>3345/2 udp</code> means UDP ports 3345 and 3346. Usually RTP uses even-numbered ports for data with the corresponding one-higher odd ports used for the RTCP session belonging to the RTP session. So <code>2910/2 RTP/AVP</code> means ports 2910 and 2912 are used for RTP and 2911 and 2913 are used for RTCP.</p> <p>Media types include <code>udp</code> for an unspecified protocol that uses UDP, <code>RTP</code> or <code>RTP/AVP</code> for standard RTP and <code>RTP/SAVP</code> for secure RTP.</p>
o=	<p>The sender's username, a session identifier, a session version number, the network type (usually IN for Internet), the address type (for example, IPv4 or IPv6), and the sending device's IP address. The o= field becomes a universal identifier for this version of this session description. For example:</p> <pre>o=PhoneA 5462346 332134 IN IP4 10.31.101.20</pre>
r=	<p>Repeat times for a session. Used if a session will be repeated at one or more timed intervals. Not normally used for VoIP calls. The times can be in different formats. For example:</p> <pre>r=7d 1h 0 25h r=604800 3600 0 90000</pre>
s=	<p>Any text that describes the session or s= followed by a space. For example:</p> <pre>s=Call from inviter</pre>
t=	<p>The start and stop time of the session. Sessions with no time restrictions (most VoIP calls) have a start and stop time of 0.</p> <pre>t=0 0</pre>
v=	<p>SDP protocol version. The current SDP version is 0 so the v= field is always:</p> <pre>v=0</pre>
z=	<p>Time zone adjustments. Used for scheduling repeated sessions that span the time between changing from standard to daylight savings time.</p> <pre>z=2882844526 -1h 2898848070 0</pre>

Example SIP messages

The following example SIP INVITE request message was sent by PhoneA to PhoneB. The first nine lines are the SIP headers. The SDP profile starts with v=0 and the media part of the session profile is the last line, starting with m=.

```
INVITE sip:PhoneB@172.20.120.30 SIP/2.0
Via: SIP/2.0/UDP 10.31.101.50:5060
From: PhoneA <sip:PhoneA@10.31.101.20>
To: PhoneB <sip:PhoneB@172.20.120.30>
Call-ID: 314159@10.31.101.20
CSeq: 1 INVITE
Contact: sip:PhoneA@10.31.101.20
Content-Type: application/sdp
Content-Length: 124
v=0
o=PhoneA 5462346 332134 IN IP4 10.31.101.20
s=Let's Talk
t=0 0
c=IN IP4 10.31.101.20
m=audio 49170 RTP 0 3
```

The following example shows a possible 200 OK SIP response message in response to the previous INVITE request message. The response includes 200 OK which indicates success, followed by an echo of the original SIP INVITE request followed by PhoneB's SDP profile.

```
SIP/2.0 200 OK
Via: SIP/2.0/UDP 10.31.101.50:5060
From: PhoneA <sip:PhoneA@10.31.101.20>
To: PhoneB <sip:PhoneB@172.20.120.30>
Call-ID: 314159@10.31.101.20
CSeq: 1 INVITE
Contact: sip:PhoneB@10.31.101.30
Content-Type: application/sdp
Content-Length: 107
v=0
o=PhoneB 124333 67895 IN IP4 172.20.120.30
s=Hello!
t=0 0
c=IN IP4 172.20.120.30
m=audio 3456 RTP 0
```

SIP can support multiple media streams for a single SIP session. Each media stream will have its own c= and m= lines in the body of the message. For example, the following message includes three media streams:

```
INVITE sip:PhoneB@172.20.120.30 SIP/2.0
Via: SIP/2.0/UDP 10.31.101.20:5060
From: PhoneA <sip:PhoneA@10.31.101.20>
To: PhoneB <sip:PhoneB@172.20.120.30>
Call-ID: 314159@10.31.101.20
CSeq: 1 INVITE
Contact: sip:PhoneA@10.31.101.20
Content-Type: application/sdp
Content-Length: 124
v=0
o=PhoneA 5462346 332134 IN IP4 10.31.101.20
s=Let's Talk
t=0 0
```

```
c=IN IP4 10.31.101.20
m=audio 49170 RTP 0 3
c=IN IP4 10.31.101.20
m=audio 49172 RTP 0 3
c=IN IP4 10.31.101.20
m=audio 49174 RTP 0 3
```

The SIP session helper

The SIP session-helper is a high-performance solution that provides basic support for SIP calls passing through the FortiGate by opening SIP and RTP pinholes and by performing NAT of the addresses in SIP messages.

The SIP session helper:

- Understands SIP dialog messages.
- Keeps the states of the SIP transactions between SIP UAs and SIP servers.
- Translates SIP header and SDP information to account for NAT operations performed by the FortiGate.
- Opens up and closes dynamic SIP pinholes for SIP signaling traffic.
- Opens up and closes dynamic RTP and RTSP pinholes for RTP and RTSP media traffic.
- Provides basic SIP security as an access control device.
- Uses the intrusion protection (IPS) engine to perform basic SIP protocol checks.

SIP session helper configuration overview

By default FortiOS uses the SIP ALG for SIP traffic. If you want to use the SIP session helper you need to enter the following commands to disable the SIP ALG and to enable the SIP session helper:

```
config system settings
    set default-voip-alg-mode kernel-helper-based
    set sip-helper enable
end
```

The SIP session helper is disabled by default and must be enabled for the SIP session helper to process VoIP traffic. The SIP session help is set to listen for SIP traffic on TCP or UDP port 5060. SIP sessions using port 5060 accepted by a security policy that does not include a VoIP profile are processed by the SIP session helper.



You might want to disable the SIP session helper and the SIP ALG if you don't want the FortiGate to apply NAT or other SIP session helper features to SIP traffic. With the SIP session helper and the SIP ALG disabled, the FortiGate can still accept SIP sessions if they are allowed by a security policy, but the FortiGate will not be able to open pinholes or NAT the addresses in the SIP messages.

You can enable and disable the SIP session helper, change the TCP or UDP port that the session helper listens on for SIP traffic, and enable or disable SIP NAT tracing. If the FortiGate is operating with multiple VDOMs, each VDOM can have a different SIP session helper configuration.

To have the SIP session helper process SIP sessions you need to add a security policy that accepts SIP sessions on the configured SIP UDP or TCP ports. The security policies can have service set to ANY, or to the SIP pre-defined firewall service, or a custom firewall service. The SIP pre-defined firewall service restricts the security policy to only accepting sessions on UDP port 5060.

If NAT is enabled for security policies that accept SIP traffic, the SIP session helper translates addresses in SIP headers and in the RDP profile and opens up pinholes as required for the SIP traffic. This includes security policies that perform source NAT and security policies that contain virtual IPs that perform destination NAT and port forwarding. No special SIP configuration is required for this address translation to occur, it is all handled automatically by the SIP session helper according to the NAT configuration of the security policy that accepts the SIP session.

To use the SIP session helper you must not add a VoIP profile to the security policy. If you add a VoIP profile, SIP traffic bypasses the SIP session helper and is processed by the SIP ALG.



In most cases you would want to use the SIP ALG since the SIP session helper provides limited functionality. However, the SIP session helper is available and can be useful for high-performance solutions where a high level of SIP security is not a requirement.

Viewing, removing, and adding the SIP session helper configuration

Enter the following command to find the sip session helper entry in the session-helper list:

```
show system session-helper
.
.
.
edit 13
  set name sip
  set port 5060
  set protocol 17
next
.
.
.
```

This command output shows that the sip session helper listens on UDP port 5060 for SIP sessions.

Enter the following command to delete session-helper list entry number 13:

```
config system session-helper
  delete 13
end
```

If you want to use the SIP session helper you can verify whether it is available using the `show system session-helper` command.

If the SIP session helper has been removed from the session-helper list you can use the following command to add it back to the session helper list:

```
config system session-helper
  edit 0
    set name sip
    set port 5060
    set protocol 17
  end
```

Changing the port numbers that the SIP session helper listens on

You can use the following command to change the port number that the SIP session helper listens on for SIP traffic to 5064. The SIP session helper listens on the same port number for UDP and TCP SIP sessions. In this example, the SIP session helper is session helper 13:

```
config system session-helper
  edit 13
    set port 5064
  end
```



The `config system settings options sip-tcp-port, sip-udp-port, and sip-ssl-port` control the ports that the SIP ALG listens on for SIP sessions. See [Changing the port numbers that the SIP ALG listens on on page 3126](#).

Your FortiGate may use a different session helper number for SIP. Enter the following command to view the session helpers:


```
show system session-helper
.
.
.
edit 13
    set name sip
    set port 5060
    set protocol 17
end
.
.
.
```

Configuration example: SIP session helper in transparent mode

The figure below shows an example SIP network consisting of a FortiGate operating in transparent mode between two SIP phones. Since the FortiGate is operating in transparent mode both phones are on the same network and the FortiGate and the SIP session helper does not perform NAT. Even though the SIP session helper is not performing NAT you can use this configuration to apply SIP session helper security features to the SIP traffic.

The FortiGate requires two security policies that accept SIP packets. One to allow SIP Phone A to start a session with SIP Phone B and one to allow SIP Phone B to start a session with SIP Phone A.

SIP network with FortiGate in transparent mode



General configuration steps

The following general configuration steps are required for this SIP configuration that uses the SIP session helper. This example includes security policies that specifically allow SIP sessions using UDP port 5060 from Phone A to Phone B and from Phone B to Phone A. In most cases you would have more than two phones so would use more general security policies. Also, you can set the firewall service to ANY to allow traffic other than SIP on UDP port 5060.

This example assumes that you have entered the following command to enable using the SIP session helper:

```
config system settings
    set default-voip-alg-mode kernel-helper-based
end
```

1. Add firewall addresses for Phone A and Phone B.
2. Add a security policy that accepts SIP sessions initiated by Phone A.
3. Add a security policy that accepts SIP sessions initiated by Phone B.

Configuration steps - GUI

To add firewall addresses for the SIP phones

1. Go to **Policy & Objects > Addresses**.
2. Select **Create New > Address** to add the following addresses for Phone A and Phone B:

Category	Address
Name	Phone_A
Type	IP/Netmask
Subnet / IP Range	10.31.101.20/255.255.255.255
Interface	port1

Category	Address
Name	Phone_B
Type	IP/Netmask
Subnet / IP Range	10.31.101.30/255.255.255.255
Interface	port2

To add security policies to accept SIP sessions

1. Go to **Policy & Objects > IPv4 Policy**.
2. Select **Create New** to add a security policy.
3. Add a security policy to allow Phone A to send SIP request messages to Phone B:

Incoming Interface	port1
Outgoing Interface	port2
Source	Phone_A
Destination Address	Phone_B
Schedule	always

Service	SIP
Action	ACCEPT

4. Select **OK**.
5. Add a security policy to allow Phone B to send SIP request messages to Phone A:

Incoming Interface	port2
Outgoing Interface	port1
Source Address	Phone_B
Destination Address	Phone_A
Schedule	always
Service	SIP
Action	ACCEPT

6. Select **OK**.

Configuration steps - CLI

To add firewall addresses for Phone A and Phone B and security policies to accept SIP sessions

1. Enter the following command to add firewall addresses for Phone A and Phone B.
2. Enter the following command to add security policies to allow Phone A to send SIP request messages to Phone B and Phone B to send SIP request messages to Phone A.

```
config firewall address
  edit Phone_A
    set associated interface port1
    set type ipmask
    set subnet 10.31.101.20 255.255.255.255
  next
  edit Phone_B
    set associated interface port2
    set type ipmask
    set subnet 10.31.101.30 255.255.255.255
  end
```

```
config firewall policy
  edit 0
    set srcintf port1
    set dstintf port2
    set srcaddr Phone_A
    set dstaddr Phone_B
    set action accept
    set schedule always
    set service SIP
  next
  edit 0
    set srcintf port2
```

```
set dstintf port1
set srcaddr Phone_B
set dstaddr Phone_A
set action accept
set schedule always
set service SIP
set utm-status enable
end
```

SIP session helper diagnose commands

You can use the `diagnose sys sip` commands to display diagnostic information for the SIP session helper.

Use the following command to set the debug level for the SIP session helper. Different debug masks display different levels of detail about SIP session helper activity.

```
diagnose sys sip debug-mask <debug_mask_int>
```

Use the following command to display the current list of SIP dialogs being processed by the SIP session help. You can also use the `clear` option to delete all active SIP dialogs being processed by the SIP session helper.

```
diagnose sys sip dialog {clear | list}
```

Use the following command to display the current list of SIP NAT address mapping tables being used by the SIP session helper.

```
diagnose sys sip mapping list
```

Use the following command to display the current SIP session helper activity including information about the SIP dialogs, mappings, and other SIP session help counts. This command can be useful to get an overview of what the SIP session helper is currently doing.

```
diagnose sys sip status
```


The SIP ALG

In most cases you should use the SIP Application Layer Gateway (ALG) for processing SIP sessions. The SIP ALG provides the same basic SIP support as the SIP session helper. Additionally, the SIP ALG provides a wide range of features that protect your network from SIP attacks, apply rate limiting to SIP sessions, check the syntax of SIP and SDP content of SIP messages, and provide detailed logging and reporting of SIP activity.

By default all SIP traffic is processed by the SIP ALG. If the policy that accepts the SIP traffic includes a VoIP profile, the SIP traffic is processed by that profile. If the policy does not include a SIP profile the SIP traffic is processed by the SIP ALG using the default VoIP profile.

If a FortiGate or a VDOM has been configured to use the SIP session helper, you can change this behavior to the default configuration of using the SIP ALG with the following command:

```
config system settings
    set default-voip-alg-mode proxy-based
    set sip-helper disable
end
```



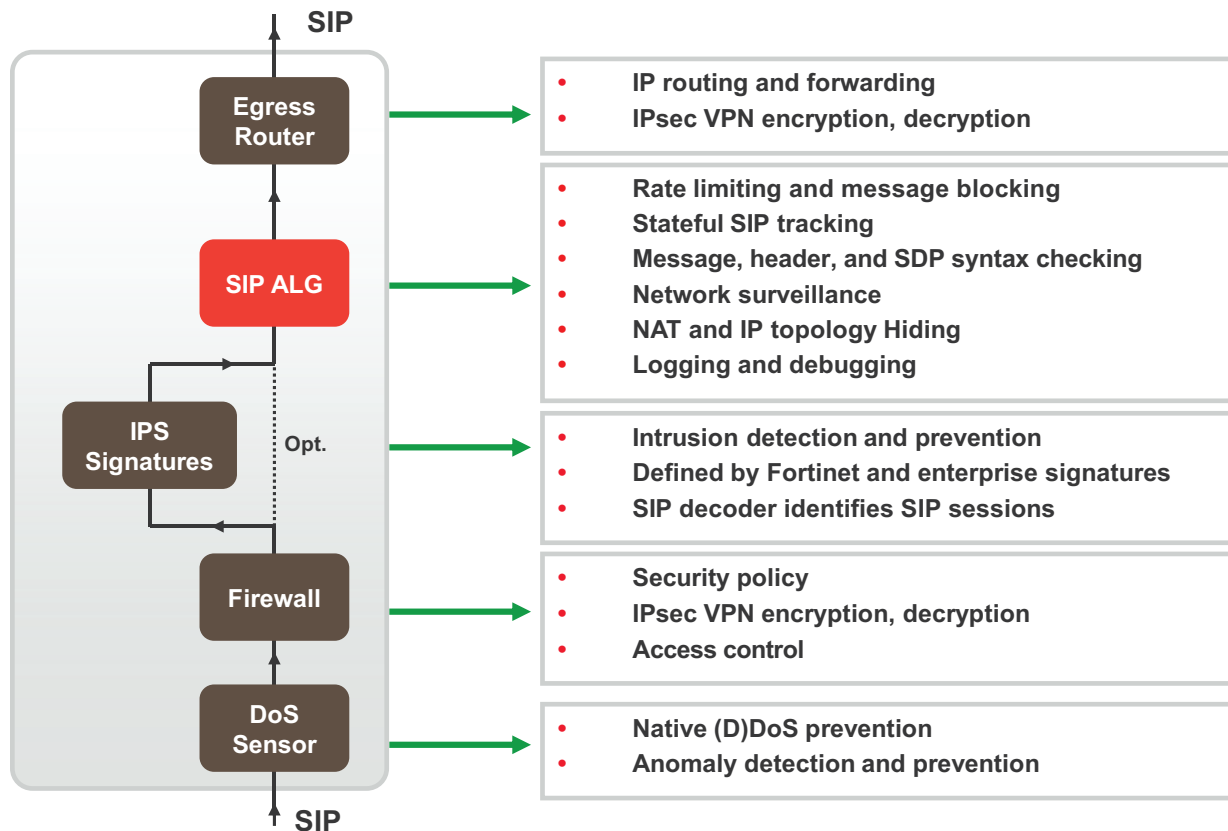
From the GUI you can only configure VoIP security profiles and add them to security policies if VoIP is turned on under **System > Feature Visibility**. However, you can always configure VoIP profiles and add them to security profiles from the CLI. And if the `default-voip-alg` mode is set to `proxy-based` the default SIP profile will still be used even if VoIP security profiles are not visible from the GUI.

As shown in the figure below, the FortiGate SIP ALG intercepts SIP packets after they have been routed by the routing module, accepted by a security policy and passed through DoS and IPS Sensors (if DoS and IPS are enabled). The ALG raises SIP packets to the application layer, analyzes the SIP and SDP addressing information in the SIP messages, makes adjustments (for example, NAT) to this addressing if required, and then sends the packets out the egress interface to their destination.

The SIP ALG provides:

- All the same features as the SIP session helper including NAT and SIP and RTP Pinholes.
- In addition for the ALG you can enable or disable RTP pinholing, SIP register pinholing and SIP contact pinholing. In a signaling only environment where the RTP stream bypasses the FortiGate, you can disable RTP pinholing to improve performance.
- SIP TCP and UDP support
- SIP Message order checking
- Configurable Header line length maximums

The SIP ALG works at the application level after ingress packets are accepted by a security policy



- Message fragment assembly (TCP)
- If SIP messages are fragmented across multiple packets, the FortiGate assembles the fragments, does inspection and pass the message in its entirety to the SIP server as one packet. This offloads the server from doing all the TCP processing of fragments.
- L4 Protocol Translation
- Message Flood Protection
- Protects a SIP server from intentional or unintentional DoS of flooding INVITE, REGISTER, and other SIP methods by allowing control of the rate that these messages pass through the FortiGate.
- SIP message type filtering
- The FortiGate can prevent specified SIP message types from passing through the FortiGate to a SIP server. For example In a voice only SIP implementation, there may be no need to permit a SUBSCRIBE message to ever make it's way to the SIP call processor. Also, if a SIP server cannot process some SIP message types you can use SIP message type filtering to block them. For example, a SIP server could have a bug that prevents it from processing certain SIP messages. In this case you can temporarily block these message types until problem with the SIP server has been fixed.
- SIP statistics and logging
- SIP over IPv6
- SIP over SSL/TLS

- Deep SIP message syntax checking (also called deep SIP header inspection or SIP fuzzing protection). Prevents attacks that use malformed SIP messages. Can check many SIP headers and SDP statements. Configurable bypass and modification options.
- Hosted NAT traversal, Resolves IP address issue in SIP and SDP lines due to NAT-PT in far end firewall. Important feature for VoIP access networks.
- SIP High Availability (HA), including active-passive clustering and session pickup (session failover) for SIP sessions.
- Geographical Redundancy. In an HA configuration, if the active SIP server fails (missing SIP heartbeat messages or SIP traffic) SIP sessions can be redirected to a secondary SIP server in another location.
- SIP per request method message rate limitation with configurable threshold for SIP message rates per request method. Protects SIP servers from SIP overload and DoS attacks.
- RTP Bypass, Supports configurations with and without RTP pinholing. May inspect and protect SIP signaling only.
- SIP NAT with IP address conservation. Performs SIP and RTP aware IP Network Address translation. Preserves the lost IP address information in the SDP profile i= line for later processing/debugging in the SIP server. See [NAT with IP address conservation on page 3159](#).
- IP topology hiding
 - The IP topology of a network can be hidden through NAT and NAPT manipulation of IP and SIP level addressing. For example, see [SIP NAT scenario: source address translation \(source NAT\) on page 3145](#).
- SIP inspection without address translation
 - The SIP ALG inspects SIP messages but addresses in the messages are not translated. This feature can be applied to a FortiGate operating in transparent mode or in NAT/Route mode. In transparent mode you add normal transparent mode security policies that enable the SIP ALG and include a VoIP profile that causes the SIP ALG to inspect SIP traffic as required. For an example configuration, see [Configuration example: SIP in transparent mode on page 3134](#).
- For a FortiGate operating in NAT/Route mode, if SIP traffic can pass between different networks without requiring NAT because is supported by the routing configuration, you can add security policies that accept SIP traffic without enabling NAT. In the VoIP profile you can configure the SIP ALG to inspect SIP traffic as required.

Enabling VoIP support from the GUI

Before you begin to configure VoIP security profiles, including SIP, from the GUI you should go to **System > Feature Visibility** and turn on **VoIP** (under **Additional Features**). VoIP settings are visible in both Inspection modes (flow and proxy).

SIP ALG configuration overview

To apply the SIP ALG, you add a SIP VoIP profile to a security policy that accepts SIP sessions. All SIP sessions accepted by the security policy will be processed by the SIP ALG using the settings in the VoIP profile. The VoIP profile contains settings that are applied to SIP, Session Initiation Protocol for Instant Messaging and Presence Leveraging Extensions (SIMPLE) and Skinny Call Control Protocol (SCCP) sessions. All SCCP sessions accepted by the security policy are also processed by the ALG. You configure SIP and SCCP settings separately. SIP settings also apply to SIMPLE sessions.

VoIP profiles

You can customize the default VoIP profile or add new VoIP profiles.

To add a new VoIP profile from the GUI go to **Security Profiles > VoIP** and select **Create New** (the + button).

For SIP, from the GUI you can configure the VoIP profile to limit the number of SIP REGISTER and INVITE requests. Many additional options for configuring how the ALG processes SIP sessions are available from the CLI.

For SCCP you can limit the call setup time. Additional SCCP options are available from the CLI.

Use the following command to add a VoIP profile named VoIP_Pro_1 from the CLI:

```
config voip profile
  edit VoIP_Pro_1
end
```

FortiGates include two pre-defined VoIP profiles. On the GUI these profiles look identical. However, the CLI-only settings result in the following functionality.

default	<p>The most commonly used VoIP profile. This profile enables both SIP and SCCP and places the minimum restrictions on what calls will be allowed to negotiate. This profile allows normal SCCP, SIP and RTP sessions and enables the following security settings:</p> <ul style="list-style-type: none"> • <code>strict-register</code> to open smaller more secure pinholes (see Enhancing SIP pinhole security on page 3165). • <code>block-long-lines</code> to block SIP messages with lines that exceed maximum line lengths. • <code>block-unknown</code> to block unrecognized SIP request messages. • <code>open-record-route-pinhole</code> to open pinholes for Record-Route messages. • <code>nat-trace</code> (see NAT with IP address conservation on page 3159). • <code>contact-fixup</code> perform NAT on the IP addresses and port numbers in SIP headers in SIP CONTACT messages even if they don't match the session's IP address and port numbers. • <code>ips-rtp</code> to enable IPS in security policies that also accept SIP sessions to protect the SIP traffic from SIP-based attacks.
strict	<p>This profile is available for users who want to validate SIP messages and to only allow SIP sessions that are compliant with RFC 3261. In addition to the settings in the default VoIP profile, the strict profile sets all SIP deep message inspection header checking options (for example, <code>malformed-request-line</code> and many others) to <code>discard</code>. So the strict profile blocks and drops SIP messages that contain malformed SIP or SDP lines that can be detected by the ALG. For more information about SIP deep header inspection, see Deep SIP message inspection on page 3175.</p>

Neither of the default profiles applies SIP rate limiting. To apply more ALG features to SIP sessions you can clone (copy) the pre-defined VoIP profiles and make your own modifications to them. You can clone VoIP profiles from the GUI or the CLI. For example, from the CLI, to clone the default profile and configure the limit for SIP NOTIFY request messages to 1000 messages per second per security policy and block SIP INFO request messages.

```
config voip profile
  clone default to my_voip_pro
  edit my_voip_pro
    config sip
      set notify-rate 1000
      set block-info enable
    end
  end
```

Changing the port numbers that the SIP ALG listens on

Most SIP configurations use TCP or UDP port 5060 for SIP sessions and port 5061 for SIP SSL sessions. If your SIP network uses different ports for SIP sessions you can use the following command to configure the SIP ALG to listen on a different TCP, UDP, or SSL ports. For example, to change the TCP port to 5064, the UDP port to 5065, and the SSL port to 5066.

```
config system settings
  set sip-tcp-port 5064
  set sip-udp-port 5065
  set sip-ssl-port 5066
end
```

You also configure the SIP ALG to listen in two different TCP ports and two different UDP ports for SIP sessions. For example, if you receive SIP TCP traffic on port 5060 and 5064 and UDP traffic on ports 5061 and 5065 you can enter the following command to receive the SIP traffic on all of these ports:

```
config system settings
  set sip-tcp-port 5060 5064
  set sip-udp-port 5061 5065
end
```

Disabling the SIP ALG in a VoIP profile

SIP is enabled by default in a VoIP profile. If you are just using the VoIP profile for SCCP you can use the following command to disable SIP in the VoIP profile.

```
config voip profile
  edit VoIP_Pro_2
    config sip
      set status disable
    end
```

SIP ALG diagnose commands

You can use the `diagnose sys sip-proxy` command to display diagnostic information for the SIP ALG. A number of options are available including:

Use the following command to list all active SIP calls being processed by the SIP ALG. You can also use the `clear` option to delete all active SIP calls being processed by the SIP ALG, the `idle` option to list idle SIP calls, and the `invite` option to list SIP invite transactions.

```
diagnose sys sip-proxy calls {clear | list | idle | invite}
```

Use the following commands to employ filters to display specific information about the SIP ALG and the session that it is processing. You can build up a filter by including a number of options such as source address, VoIP profile, policy, and so on.

```
diagnose sys sip-proxy filter <filter_options>
```

```
diagnose sys sip-proxy log-filter <filter_options>
```

Use the following command to display the active SIP rate limiting meters and their current settings.

```
diagnose sys sip-proxy meters list
```

Use the following command to display status information about the SIP sessions being processed by the SIP ALG. You can also clear all SIP ALG statistics.

```
diagnose sys sip-proxy stats {clear | list}
```

Conflicts between the SIP ALG and the session helper

If you suspect that the SIP session helper is being used instead of the ALG, you can use the `diagnose sys sip` command to determine if the SIP session helper is processing SIP sessions. For example, the following command displays the overall status of the SIP sessions being processed by the SIP session helper:



The `diagnose sys sip` command only displays current status information. To see activity the SIP session helper has to actually be processing SIP sessions when you enter the command. For example, if the SIP session helper had been used for processing calls that ended 5 minutes ago, the command output would show no SIP session helper activity.

```
diagnose sys sip status
dialogs: max=32768, used=0
mappings: used=0
dialog hash by ID: size=2048, used=0, depth=0
dialog hash by RTP: size=2048, used=0, depth=0
mapping hash: size=2048, used=0, depth=0
count0: 0
count1: 0
count2: 0
count3: 0
count4: 0
```

This command output shows that the session helper is not processing SIP sessions because all of the used and count fields are 0. If any of these fields contains non-zero values then the SIP session helper may be processing SIP sessions.

Also, you can check to see if some ALG-only features are not being applied to all SIP sessions. For example, FortiView pages displays statistics for SIP and SCCP calls processed by the ALG but not for calls processed by the session helper. So if you see fewer calls than expected the session helper may be processing some of them.

Finally, you can check the policy usage and session information dashboard widgets to see if SIP sessions are being accepted by the wrong security policies.

Stateful SIP tracking, call termination, and session inactivity timeout

The SIP ALG tracks SIP dialogs over their lifespan between the first INVITE message and the Final 200 OK and ACK messages. For every SIP dialog, stateful SIP tracking reviews every SIP message and makes adjustment to SIP tracking tables as required. These adjustments include source and destination IP addresses, address translation, dialog expiration information, and media stream port changes. Such changes can also result in dynamically opening and closing pinholes. You can use the `diagnose sys sip-proxy stats list` and the `diagnose sys sip-proxy filter` command to view the SIP call data being tracked by the SIP ALG.

The SIP ALG uses the SIP Expires header line to time out a SIP dialog if the dialog is idle and a Re-INVITE or UPDATE message is not received. The SIP ALG gets the Session-Expires value, if present, from the 200 OK response to the INVITE message. If the SIP ALG receives an INVITE before the session times out, all timeout values are reset to the settings in the new INVITE message or to default values. As a precautionary measure, the SIP ALG uses hard timeout values to set the maximum amount of time a call can exist. This ensures that the FortiGate is protected if a call ends prematurely.

When a SIP dialog ends normally, the SIP ALG deletes the SIP call information and closes open pinholes. A SIP call can also end abnormally due to an unexpected signaling or transport event that cuts off the call. When a call ends abnormally the SIP messages to end the call may not be sent or received. A call can end abnormally for the following reasons:

- Phones or servers crash during a call and a BYE message is not received.
- To attack a SIP system, a malicious user never send a BYE message.
- Poor implementations of SIP fail to process Record-Route messages and never send a BYE message.
- Network failures prevent a BYE message from being received.

Any phone or server in a SIP call can cancel the call by sending a CANCEL message. When a CANCEL message is received by the FortiGate, the SIP ALG closes open pinholes. Before terminating the call, the ALG waits for the final 200 OK message.

The SIP ALG can be configured to terminate SIP calls if the SIP dialog message flow or the call RTP (media) stream is interrupted and does not recover. You can use the following commands to configure terminating inactive SIP sessions and to set timers or counters to control when the call is terminated by the SIP ALG.

Adding a media stream timeout for SIP calls

Use the following command in a VoIP profile to terminate SIP calls accepted by a security policy containing the VoIP profile when the RTP media stream is idle for 100 seconds.

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set call-keepalive 100
    end
  end
```

You can adjust this setting between 1 and 10,080 seconds. The default call keepalive setting of 0 disables terminating a call if the media stream is interrupted. Set call keepalive higher if your network has latency problems that could temporarily interrupt media streams. If you have configured call keepalive and the FortiGate terminates calls unexpectedly you can increase the call keepalive time to resolve the problem.



Call keep alive should be used with caution because enabling this feature results in extra FortiGate CPU overhead and can cause delay/jitter for the VoIP call. Also, the FortiGate terminates the call without sending SIP messages to end the call. And if the SIP endpoints send SIP messages to terminate the call they will be blocked by the FortiGate if they are sent after the FortiGate terminates the call.

Adding an idle dialog setting for SIP calls

Use the following command in a VoIP profile to terminate SIP calls when for a single security policy, when the configured number of SIP calls (or dialogs) has stopped receiving SIP messages or has not received legitimate SIP messages. Using this command you can configure how many dialogs that have been accepted by a security policy that the VoIP profile is added to become idle before the SIP ALG deletes the oldest ones. The following command sets the maximum number of idle dialogs to 200:

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set max-idle-dialogs 200
    end
  end
```

Idle dialogs would usually be dialogs that have been interrupted because of errors or problems or as the result of a SIP attack that opens a large number of SIP dialogs without closing them. This command provides a way to remove these dialogs from the dialog table and recover memory and resources being used by these open and idle dialogs.

You can adjust this setting between 1 and a very high number. The default maximum idle dialogs setting of 0 disables this feature. Set maximum dialogs higher if your network has latency problems that could temporarily interrupt SIP messaging. If you have configured max idle dialogs and the FortiGate terminates calls unexpectedly you can increase the max idle dialogs number to resolve the problem.

Changing how long to wait for call setup to complete

In some cases and some configurations your SIP system may experience delays during call setup. If this happens, some SIP ALG timers may expire before call setup is complete and drop the call. In some cases you may also want to reduce the amount of time the SIP ALG allows for call setup to complete.

You can use the `provisional-invite-expiry-time` SIP VoIP profile option to control how long the SIP ALG waits for provisional INVITE messages before assuming that the call setup has been interrupted and the SIP call should be dropped. The default value for this timer is 210 seconds. You can change it to between 10 and 3600 seconds.

Use the following command to change the expiry time to 100 seconds.

```
config voip profile
  edit Profile_name
    config sip
      set provisional-invite-expiry-time 100
    end
  end
```

SIP and RTP/RTCP

FortiGates support the Real Time Protocol (RTP) application layer protocol for the VoIP call audio stream. RTP uses dynamically assigned port numbers that can change during a call. SIP control messages that start a call and that are sent during the call inform callers of the port number to use and of port number changes during the call.

During a call, each RTP session will usually have a corresponding Real Time Control Protocol (RTCP) session. By default, the RTCP session port number is one higher than the RTP port number.

The RTP port number is included in the `m=` part of the SDP profile. In the example above, the SIP INVITE message includes RTP port number is 49170 so the RTCP port number would be 49171. In the SIP response message the RTP port number is 3456 so the RTCP port number would be 3457.

How the SIP ALG creates RTP pinholes

The SIP ALG requires the following information to create a pinhole. The SIP ALG finds this information in SIP messages and some is provided by the SIP ALG:

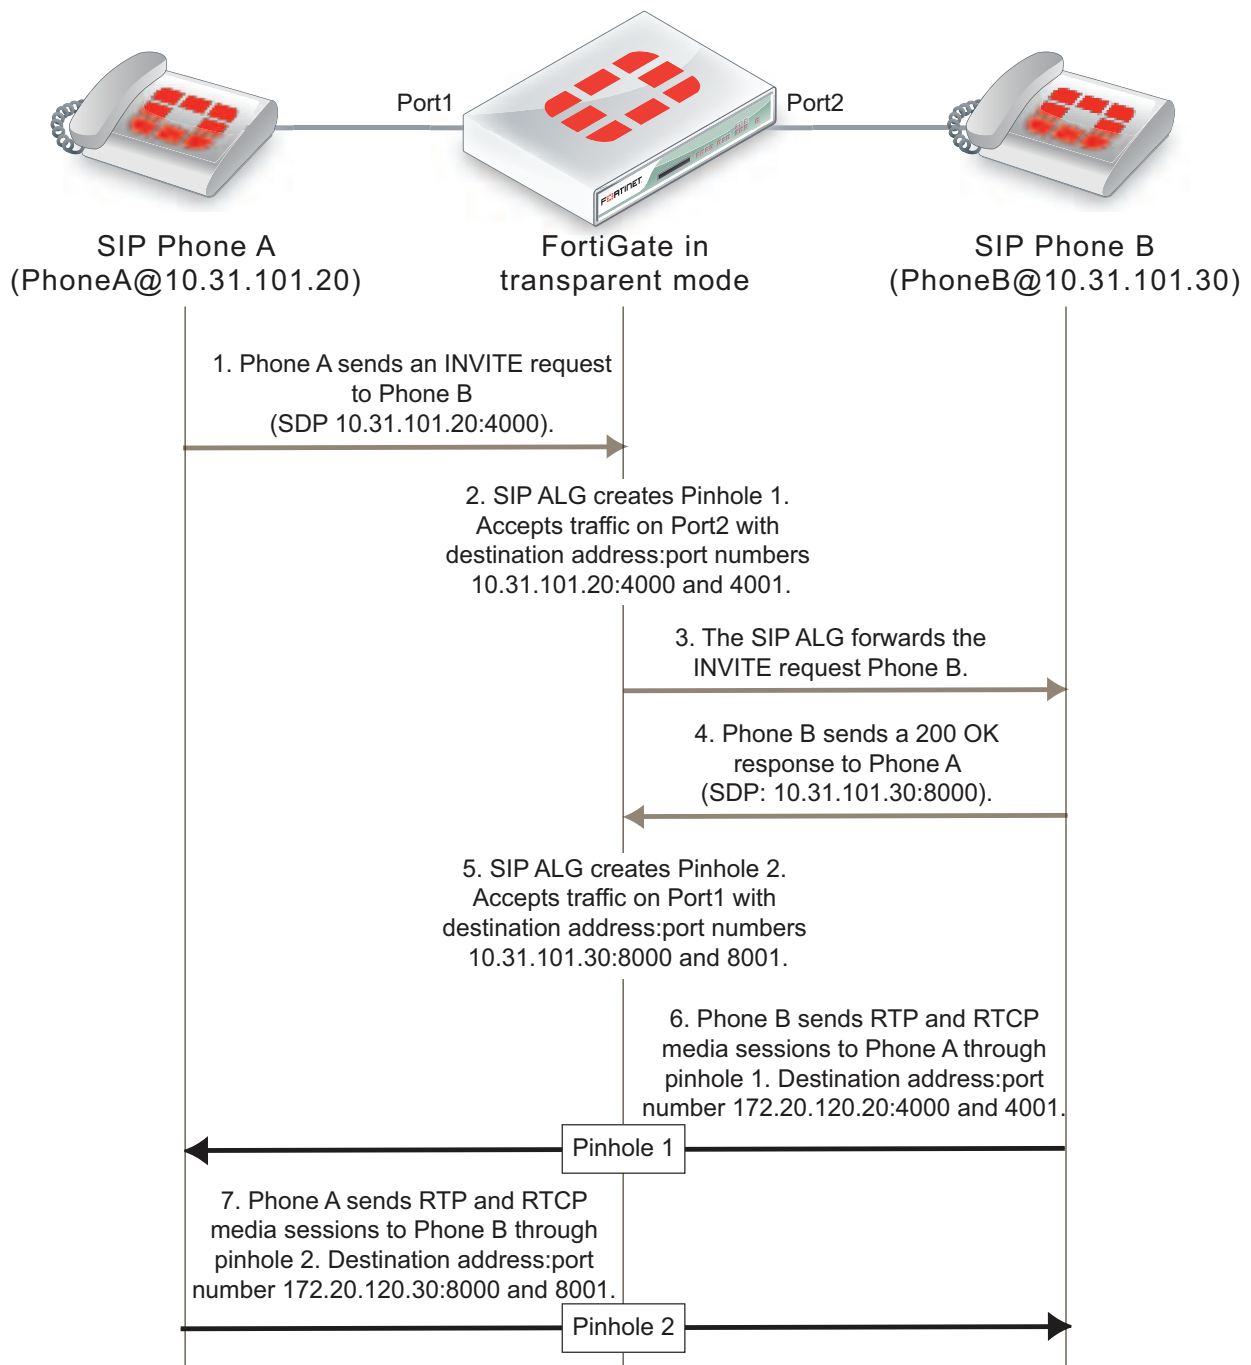
Protocol	UDP (Extracted from SIP messages by the SIP ALG.)
Source IP	Any
Source port	Any
Destination IP	The SIP ALG extracts the destination IP address from the c= line in the SDP profile. The c= line can appear in either the session or media part of the SDP profile. The SIP ALG uses the IP address in the c= line of the media part of the SDP profile first. If the media part does not contain a c= line, the SIP ALG checks the c= line in the session part of the SDP profile. If the session part of the profile doesn't contain a c= line the packet is dropped. Pinholes for RTP and RTCP sessions share the same destination IP address.
Destination port	The SIP ALG extracts the destination port number for RTP from the m= field and adds 1 to this number to get the RTCP port number.
Lifetime	The length of time during which the pinhole will be open. When the lifetime ends, the SIP ALG removes the pinhole.

The SIP ALG keeps RTP pinholes open as long as the SIP session is alive. When the associated SIP session is terminated by the SIP ALG or the SIP phones or servers participating in the call, the RTP pinhole is closed.

The figure below shows a simplified call setup sequence that shows how the SIP ALG opens pinholes. Phone A and Phone B are installed on either side of a FortiGate operating in transparent mode. Phone A and Phone B are on the same subnet. The FortiGate includes a security policy that accepts SIP sessions from port1 to port2 and from port2 to port1. The FortiGate does not require an RTP security policy, just the SIP policy.

You can see from this diagram that the SDP profile in the INVITE request from Phone A indicates that Phone A is expecting to receive a media stream sent to its IP address using port 4000 for RTP and port 4001 for RTCP. The SIP ALG creates pinhole 1 to allow this media traffic to pass through the FortiGate. Pinhole 1 is opened on the Port2 interface and will accept media traffic sent from Phone B to Phone A.

When Phone B receives the INVITE request from Phone A, Phone B will know to send media streams to Phone A using destination IP address 10.31.101.20 and ports 4000 and 4001. The 200 OK response sent from Phone B indicates that Phone B is expecting to receive a media stream sent to its IP address using ports 8000 and 8001. The SIP ALG creates pinhole 2 to allow this media traffic to pass through the FortiGate. Pinhole 2 is opened on the Port1 interface and will accept media traffic sent from Phone A to Phone B.


SIP call setup with a FortiGate in transparent mode

Configuration example: SIP in transparent mode

The figure below shows an example SIP network consisting of a FortiGate operating in transparent mode between two SIP phones. Since the FortiGate is operating in transparent mode both phones are on the same network and the FortiGate and the SIP ALG does not perform NAT. Even though the SIP ALG is not performing NAT you can use this configuration to apply SIP security features to the SIP traffic.

The FortiGate requires two security policies that accept SIP packets. One to allow SIP Phone A to start a session with SIP Phone B and one to allow SIP Phone B to start a session with SIP Phone A.

SIP network with FortiGate in transparent mode



General configuration steps

The following general configuration steps are required for this SIP configuration. This example uses the default VoIP profile. The example also includes security policies that specifically allow SIP sessions using UDP port 5060 from Phone A to Phone B and from Phone B to Phone A. In most cases you would have more than two phones so would use more general security policies. Also, you can set the security service to ANY to allow traffic other than SIP on UDP port 5060.

1. Add firewall addresses for Phone A and Phone B.
2. Add a security policy that accepts SIP sessions initiated by Phone A and includes the default VoIP profile.
3. Add a security policy that accepts SIP sessions initiated by Phone B and includes the default VoIP profile.

Configuration steps - GUI



Before you begin this procedure you may have to go to **System > Feature Visibility** and turn on VoIP.

To add firewall addresses for the SIP phones

1. Go to **Policy & Objects > Addresses**.
2. Add the following addresses for Phone A and Phone B:

Category	Address
Name	Phone_A
Type	IP/Netmask
Subnet / IP Range	10.31.101.20/255.255.255.255
Interface	port1

Category	Address
Name	Phone_B
Type	IP/Netmask
Subnet / IP Range	10.31.101.30/255.255.255.255
Interface	port2

To add security policies to apply the SIP ALG to SIP sessions

1. Go to **Policy & Objects > IPv4 Policy**.
2. Add a security policy to allow Phone A to send SIP request messages to Phone B:

Incoming Interface	port1
Outgoing Interface	port2
Source	Phone_A
Destination Address	Phone_B
Schedule	always
Service	SIP
Action	ACCEPT

3. Turn on **VoIP** and select the **default** VoIP profile.
4. Select **OK**.
5. Add a security policy to allow Phone B to send SIP request messages to Phone A:

Incoming Interface	port2
Outgoing Interface	port1
Source	Phone_B
Destination Address	Phone_A

Schedule	always
Service	SIP
Action	ACCEPT

6. Turn on **VoIP** and select the **default** VoIP profile.
7. Select **OK**.

Configuration steps - CLI

To add firewall addresses for Phone A and Phone B and security policies to apply the SIP ALG to SIP sessions

1. Enter the following command to add firewall addresses for Phone A and Phone B.

```
config firewall address
edit Phone_A
set associated-interface port1
set type ipmask
set subnet 10.31.101.20 255.255.255.255
next
edit Phone_B
set associated-interface port2
set type ipmask
set subnet 10.31.101.30 255.255.255.255
end
```

2. Enter the following command to add security policies to allow Phone A to send SIP request messages to Phone B and Phone B to send SIP request messages to Phone A.

```
config firewall policy
edit 0
set srcintf port1
set dstintf port2
set srcaddr Phone_A
set dstaddr Phone_B
set action accept
set schedule always
set service SIP
set utm-status enable
set voip-profile default
next
edit 0
set srcintf port2
set dstintf port1
set srcaddr Phone_B
set dstaddr Phone_A
set action accept
set schedule always
set service SIP
set utm-status enable
set voip-profile default
end
```

RTP enable/disable (RTP bypass)

You can configure the SIP ALG to stop from opening RTP pinholes. Called RTP bypass, this configuration can be used when you want to apply SIP ALG features to SIP signaling messages but do not want the RTP media streams to pass through the FortiGate. The FortiGate only acts as a signaling firewall and RTP media session bypass the FortiGate and no pinholes need to be created.

Enter the following command to enable RTP bypass in a VoIP profile by disabling opening RTP pinholes:

```
config voip profile
  edit VoIP_Pro_1
    config sip
      set rtp disable
    end
  end
```


Opening and closing SIP register, contact, via and record-route pinholes

You can use the `open-register-pinhole`, `open-contact-pinhole`, `open-via-port`, and `open-record-route-pinhole` VoIP profile CLI options to control whether the FortiGate opens various pinholes.

If `open-register-pinhole` is enabled (the default setting) the FortiGate opens pinholes for SIP Register request messages. You can disable `open-register-pinhole` so that the FortiGate does not open pinholes for SIP Register request messages.

If `open-contact-pinhole` is enabled (the default setting) the FortiGate opens pinholes for non-Register SIP request messages. You can disable `open-contact-pinhole` so that the FortiGate does not open pinholes for non-register requests. Non-register pinholes are usually opened for SIP INVITE requests.

If `open-via-pinhole` is disabled (the default setting) the FortiGate does not open pinholes for Via messages. You can enable `open-via-pinhole` so that the FortiGate opens pinholes for Via messages.

If `open-record-route-pinhole` is enabled (the default setting) the FortiGate opens pinholes for Record-Route messages. You can disable `open-record-route-pinhole` so that the FortiGate does not open pinholes for Record-Route messages.

Usually you would want to open these pinholes. Keeping them closed may prevent SIP from functioning properly through the FortiGate. They can be disabled, however, for interconnect scenarios (where all SIP traffic is between proxies and traveling over a single session). In some cases these settings can also be disabled in access scenarios if it is known that all users will be registering regularly so that their contact information can be learned from the register request.

You might want to prevent pinholes from being opened to avoid creating a pinhole for every register or non-register request. Each pinhole uses additional system memory, which can affect system performance if there are hundreds or thousands of users, and requires refreshing which can take a relatively long amount of time if there are thousands of active calls.

To configure a VoIP profile to prevent opening register and non-register pinholes:

```
config voip profile
  edit VoIP_Pro_1
    config sip
      set open-register-pinhole disable
      set open-contact-pinhole disable
    end
  end
```

In some cases you may not want to open pinholes for the port numbers specified in SIP Contact headers. For example, in an interconnect scenario when a FortiGate is installed between two SIP servers and the only SIP traffic through the FortiGate is between these SIP servers pinholes may not need to be opened for the port numbers specified in the Contact header lines.

If you disable `open-register-pinhole` then pinholes are not opened for ports in Contact header lines in SIP Register messages. If you disable `open-contact-pinhole` then pinholes are not opened for ports in Contact header lines in all SIP messages except SIP Register messages.

Accepting SIP register responses

You can enable the VoIP profile `open-via-pinhole` options to accept a SIP Register response message from a SIP server even if the source port of the Register response message is different from the destination port.

Most SIP servers use 5060 as the source port in the SIP register response. Some SIP servers, however, may use a different source port. If your SIP server uses a different source port, you can enable `open-via-pinhole` and the SIP ALG will create a temporary pinhole when the Register request from a SIP client includes a different source port. The FortiGate will accept a SIP Register response with any source port number from the SIP server.

Enter the following command to enable accepting any source port from a SIP server:

```
config voip profile
  edit VoIP_Pro_1
    config sip
      set open-via-pinhole enable
    end
  end
end
```

How the SIP ALG performs NAT

In most Network Address Translation (NAT) configurations, multiple hosts in a private network share a single public IP address to access the Internet. For sessions originating on the private network for the Internet, NAT replaces the private IP address of the PC in the private subnet with the public IP address of the NAT device. The NAT device converts the public IP address for responses from the Internet back into the private address before sending the response over the private network to the originator of the session.

Using NAT with SIP is more complex because of the IP addresses and media stream port numbers used in SIP message headers and bodies. When a caller on the private network sends a SIP message to a phone or SIP server on the Internet, the SIP ALG must translate the private network addresses in the SIP message to IP addresses and port numbers that are valid on the Internet. When the response message is sent back to the caller, the SIP ALG must translate these addresses back to valid private network addresses.

In addition, the media streams generated by the SIP session are independent of the SIP message sessions and use varying port numbers that can also change during the media session. The SIP ALG opens pinholes to accept these media sessions, using the information in the SIP messages to determine the pinholes to open. The ALG may also perform port translation on the media sessions.

When an INVITE message is received by the SIP ALG, the FortiGate extracts addressing and port number information from the message header and stores it in a SIP dialog table. Similar to an IP session table the data in the dialog table is used to translate addresses in subsequent SIP messages that are part of the same SIP call.

When the SIP ALG receives a response to the INVITE message arrives, (for example, an ACK or 200 OK), the SIP ALG compares the addresses in the message fields against the entries in the SIP dialog table to identify the call context of the message. The SIP ALG then translates addresses in the SIP message before forwarding them to their destination.

The addressing and port number information in SDP fields is used by the ALG to reserve ports for the media session and create a NAT mapping between them and the ports in the SDP fields. Because SDP uses sequential ports for the RTP and RTCP channels, the ALG provides consecutive even-odd ports.

SIP ALG source address translation

When a SIP call is started by a phone on a private network destined for a phone on the Internet, only source address translation is required. The phone on the private network attempts to contact the actual IP address of the phone on the Internet. However, the source address of the phone on the private network is not routable on the Internet so the SIP ALG must translate all private IP addresses in the SIP message into public IP addresses.

To configure the FortiGate for source address translation you add security policy that accepts sessions from the internal network destined for the Internet. You must enable NAT for the security policy and add a VoIP profile.

When a SIP request is received from the internal to the external network, the SIP ALG replaces the private network IP addresses and port numbers in the SIP message with the IP address of the FortiGate interface connected to the Internet. Depending on the content of the message, the ALG translates addresses in the Via:, Contact:, Route:, and Record-Route: SIP header fields. The message is then forwarded to the destination (either a VoIP phone or a SIP server on the Internet).

The VoIP phone or server in the Internet sends responses to these SIP messages to the external interface of the FortiGate. The addresses in the response messages are translated back into private network addresses and the response is forwarded to the originator of the request.

For the RTP communication between the SIP phones, the SIP ALG opens pinholes to allow media through the FortiGate on the dynamically assigned ports negotiated based on information in the SDP and the Via:, Contact:, and Record-Route: header fields. The pinholes also allow incoming packets to reach the Contact:, Via:, and Record-Route: IP addresses and ports. When processing return traffic, the SIP ALG inserts the original Contact:, Via:, Route:, and Record-Route: SIP fields back into the packets.

SIP ALG destination address translation

Incoming calls are directed from a SIP phone on the Internet to the interface of the FortiGate connected to the Internet. To receive these calls you must add a security policy to accept SIP sessions from the Internet. The security policy requires a firewall virtual IP. SIP INVITE messages from the Internet connect to the external IP address of the virtual IP. The SIP ALG uses the destination address translation defined in the virtual IP to translated the addresses in the SIP message to addresses on the private network.

When a 200 OK response message arrives from the private network, the SIP ALG translates the addresses in the message to Internet addresses and opens pinholes for media sessions from the private network to the Internet.

When the ACK message is received for the 200 OK, it is also intercepted by the SIP ALG. If the ACK message contains SDP information, the SIP ALG checks to determine if the IP addresses and port numbers are not changed from the previous INVITE. If they are, the SIP ALG deletes pinholes and creates new ones as required. The ALG also monitors the Via:, Contact:, and Record-Route: SIP fields and opens new pinholes as required.

SIP call re-invite messages

SIP Re-INVITE messages can dynamically add and remove media sessions during a call. When new media sessions are added to a call the SIP ALG opens new pinholes and update SIP dialog data. When media sessions are ended, the SIP ALG closes pinholes that are no longer needed and removes SIP dialog data.

How the SIP ALG translates IP addresses in SIP headers

The SIP ALG applies NAT to SIP sessions by translating the IP addresses contained in SIP headers. For example, the following SIP message contains most of the SIP fields that contain addresses that need to be translated:

```
INVITE PhoneB@172.20.120.30 SIP/2.0
Via: SIP/2.0/UDP 172.20.120.50:5434
From: PhoneA@10.31.101.20
To: PhoneB@172.20.120.30
Call-ID: a12abcde@172.20.120.50
Contact: PhoneA@10.31.101.20:5434
Route: <sip:example@172.20.120.50:5060>
Record-Route: <sip:example@172.20.120.50:5060>
```

How IP address translation is performed depends on whether source NAT or destination NAT is applied to the session containing the message:

Source NAT translation of IP addresses in SIP messages

Source NAT translation occurs for SIP messages sent from a phone or server on a private network to a phone or server on the Internet. The source addresses in the SIP header fields of the message are typically set to IP addresses on the private network. The SIP ALG translates these addresses to the address the FortiGate interface connected to the Internet.

Source NAT translation of IP addresses in SIP request messages

SIP header	NAT action
To:	None
From:	Replace private network address with IP address of FortiGate interface connected to the Internet.
Call-ID:	Replace private network address with IP address of FortiGate interface connected to the Internet.
Via:	Replace private network address with IP address of FortiGate interface connected to the Internet.
Request-URI:	None
Contact:	Replace private network address with IP address of FortiGate interface connected to the Internet.
Record-Route:	Replace private network address with IP address of FortiGate interface connected to the Internet.
Route:	Replace private network address with IP address of FortiGate interface connected to the Internet.

Response messages from phones or servers on the Internet are sent to the FortiGate interface connected to the Internet where the destination addresses are translated back to addresses on the private network before forwarding the SIP response message to the private network.

Source NAT translation of IP addresses in SIP response messages

SIP header	NAT action
To:	None
From:	Replace IP address of FortiGate interface connected to the Internet with private network address.
Call-ID:	Replace IP address of FortiGate interface connected to the Internet with private network address.
Via:	Replace IP address of FortiGate interface connected to the Internet with private network address.
Request-URI:	N/A
Contact:	None

SIP header	NAT action
Record-Route:	Replace IP address of FortiGate interface connected to the Internet with private network address.
Route:	Replace IP address of FortiGate interface connected to the Internet with private network address.

Destination NAT translation of IP addresses in SIP messages

Destination NAT translation occurs for SIP messages sent from a phone or server on the Internet to a firewall virtual IP address. The destination addresses in the SIP header fields of the message are typically set to the virtual IP address. The SIP ALG translates these addresses to the address of a SIP server or phone on the private network on the other side of the FortiGate.

Destination NAT translation of IP addresses in SIP request messages

SIP header	NAT action
To:	Replace VIP address with address on the private network as defined in the firewall virtual IP.
From:	None
Call-ID:	None
Via:	None
Request-URI:	Replace VIP address with address on the private network as defined in the firewall virtual IP.
Contact:	None
Record-Route:	None
Route:	None

SIP response messages sent in response to the destination NAT translated messages are sent from a server or a phone on the private network back to the originator of the request messages on the Internet. These reply messages are accepted by the same security policy that accepted the initial request messages. The firewall VIP in the original security policy contains the information that the SIP ALG uses to translate the private network source addresses in the SIP headers into the firewall virtual IP address.

Destination NAT translation of IP addresses in SIP response messages

SIP header	NAT action
To:	None

SIP header	NAT action
From:	Replace private network address with firewall VIP address.
Call-ID:	None
Via:	None
Request-URI:	N/A
Contact:	Replace private network address with firewall VIP address.
Record-Route:	Replace private network address with firewall VIP address.
Route:	None

How the SIP ALG translates IP addresses in the SIP body

The SDP session profile attributes in the SIP body include IP addresses and port numbers that the SIP ALG uses to create pinholes for the media stream.

The SIP ALG translates IP addresses and port numbers in the `o=`, `c=`, and `m=` SDP lines. For example, in the following lines the ALG could translate the IP addresses in the `o=` and `c=` lines and the port number (49170) in the `m=` line.

```
o=PhoneA 5462346 332134 IN IP4 10.31.101.20
c=IN IP4 10.31.101.20
m=audio 49170 RTP 0 3
```

If the SDP session profile includes multiple RTP media streams, the SIP ALG opens pinholes and performs the required address translation for each one.

The two most important SDP attributes for the SIP ALG are `c=` and `m=`. The `c=` attribute is the connection information attribute. This field can appear at the session or media level. The syntax of the connection attribute is:

```
c=IN {IPv4 | IPv6} <destination_ip_address>
```

Where

- `IN` is the network type. FortiGates support the `IN` or Internet network type.
- `{IPv4 | IPv6}` is the address type. FortiGates support IPv4 or IPv6 addresses in SDP statements. However, FortiGates do not support all types of IPv6 address translation. See [SIP over IPv6 on page 3174](#).
- `<destination_IP_address>` is the unicast numeric destination IP address or domain name of the connection in either IPv4 or IPv6 format.

The syntax of the media attribute is:

```
m=audio <port_number> RTP <format_list>
```

Where

- `audio` is the media type. FortiGates support the `audio` media type.
- `<port_number>` is the destination port number used by the media stream.
- `RTP` is the application layer transport protocol used for the media stream. FortiGates support the Real Time Protocol (RTP) transport protocol.

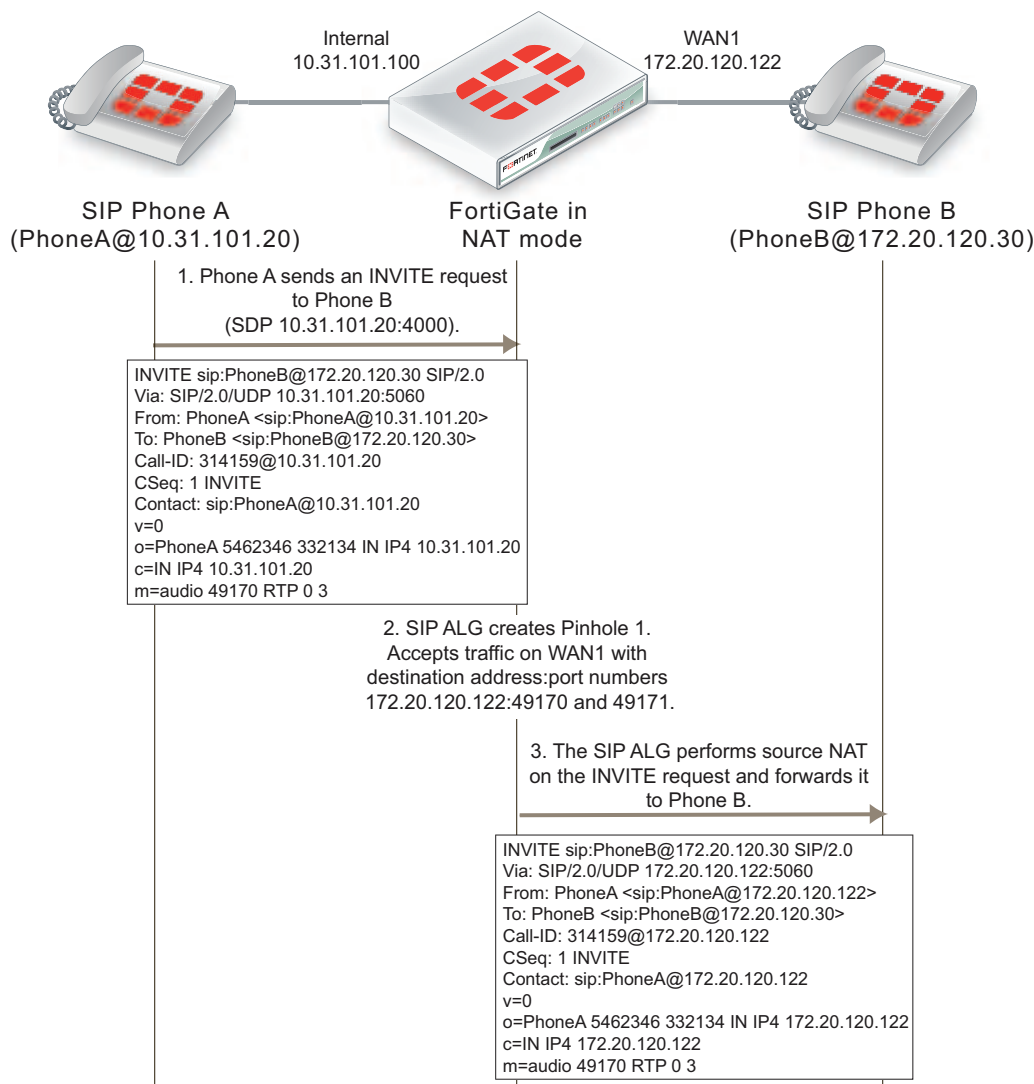
- `<format_list>` is the format list that provides information about the application layer protocol that the media uses.

SIP NAT scenario: source address translation (source NAT)

The following figures show a source address translation scenario involving two SIP phones on different networks, separated by a FortiGate. In the scenario, SIP Phone A sends an INVITE request to SIP Phone B and SIP Phone B replies with a 200 OK response and then the two phones start media streams with each other.

To simplify the diagrams, some SIP messages are not included (for example, the Ringing and ACK response messages) and some SIP header lines and SDP profile lines have been removed from the SIP messages.

SIP source NAT scenario part 1: INVITE request sent from Phone A to Phone B

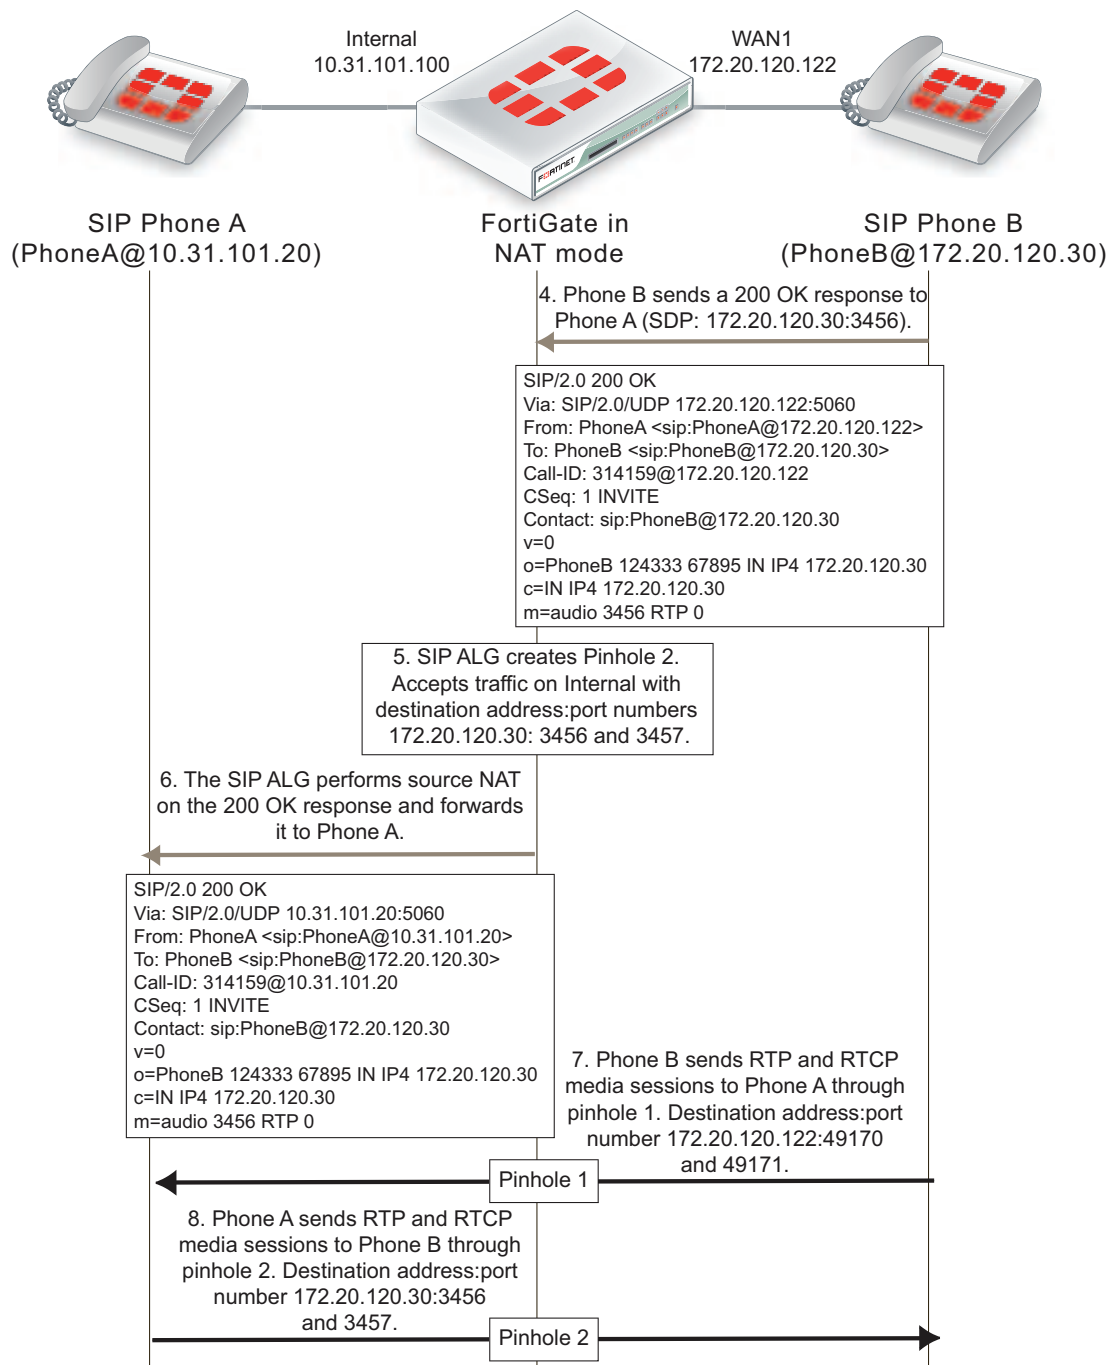


For the replies to SIP packets sent by Phone A to be routable on Phone B's network, the FortiGate uses source NAT to change their source address to the address of the WAN1 interface. The SIP ALG makes similar changes to the source addresses in the SIP headers and SDP profile. For example, the original INVITE request from Phone A includes the address of Phone A (10.31.101.20) in the from header line. After the INVITE request passes through

the FortiGate, the address of Phone A in the From SIP header line is translated to 172.20.120.122, the address of the FortiGate WAN1 interface. As a result, Phone B will reply to SIP messages from Phone A using the WAN1 interface IP address.

The FortiGate also opens a pinhole so that it can accept media sessions sent to the WAN1 IP address using the port number in the m= line of the INVITE request and forward them to Phone A after translating the destination address to the IP address of Phone A.

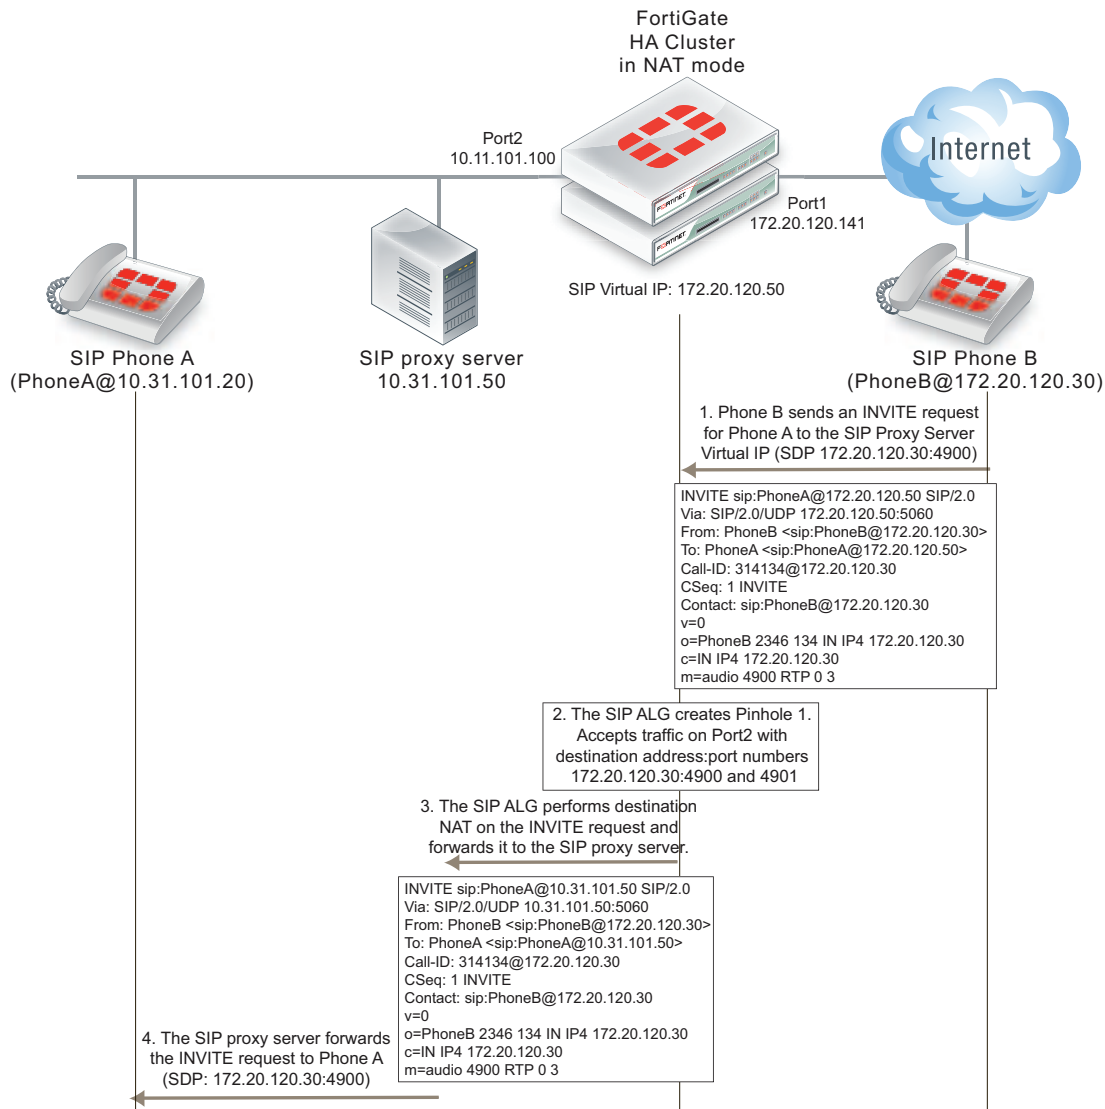
Phone B sends the 200 OK response to the INVITE message to the WAN1 interface. The SDP profile includes the port number that Phone B wants to use for its media stream. The FortiGate forwards 200 OK response to Phone A after translating the addresses in the SIP and SDP lines back to the IP address of Phone A. The SIP ALG also opens a pinhole on the Internal interface that accepts media stream sessions from Phone A with destination address set to the IP address of Phone B and using the port that Phone B added to the SDP m= line.

SIP source NAT scenario part 2: 200 OK returned and media streams established**SIP NAT scenario: destination address translation (destination NAT)**

The following figures show how the SIP ALG translates addresses in a SIP INVITE message sent from SIP Phone B on the Internet to SIP Phone A on a private network using the SIP proxy server. Because the addresses on the private network are not visible from the Internet, the security policy on the FortiGate that accepts SIP sessions

includes a virtual IP. Phone A sends SIP the INVITE message to the virtual IP address. The FortiGate accepts the INVITE message packets and using the virtual IP, translates the destination address of the packet to the IP address of the SIP proxy server and forwards the SIP message to it.

SIP destination NAT scenario part 1: INVITE request sent from Phone B to Phone A



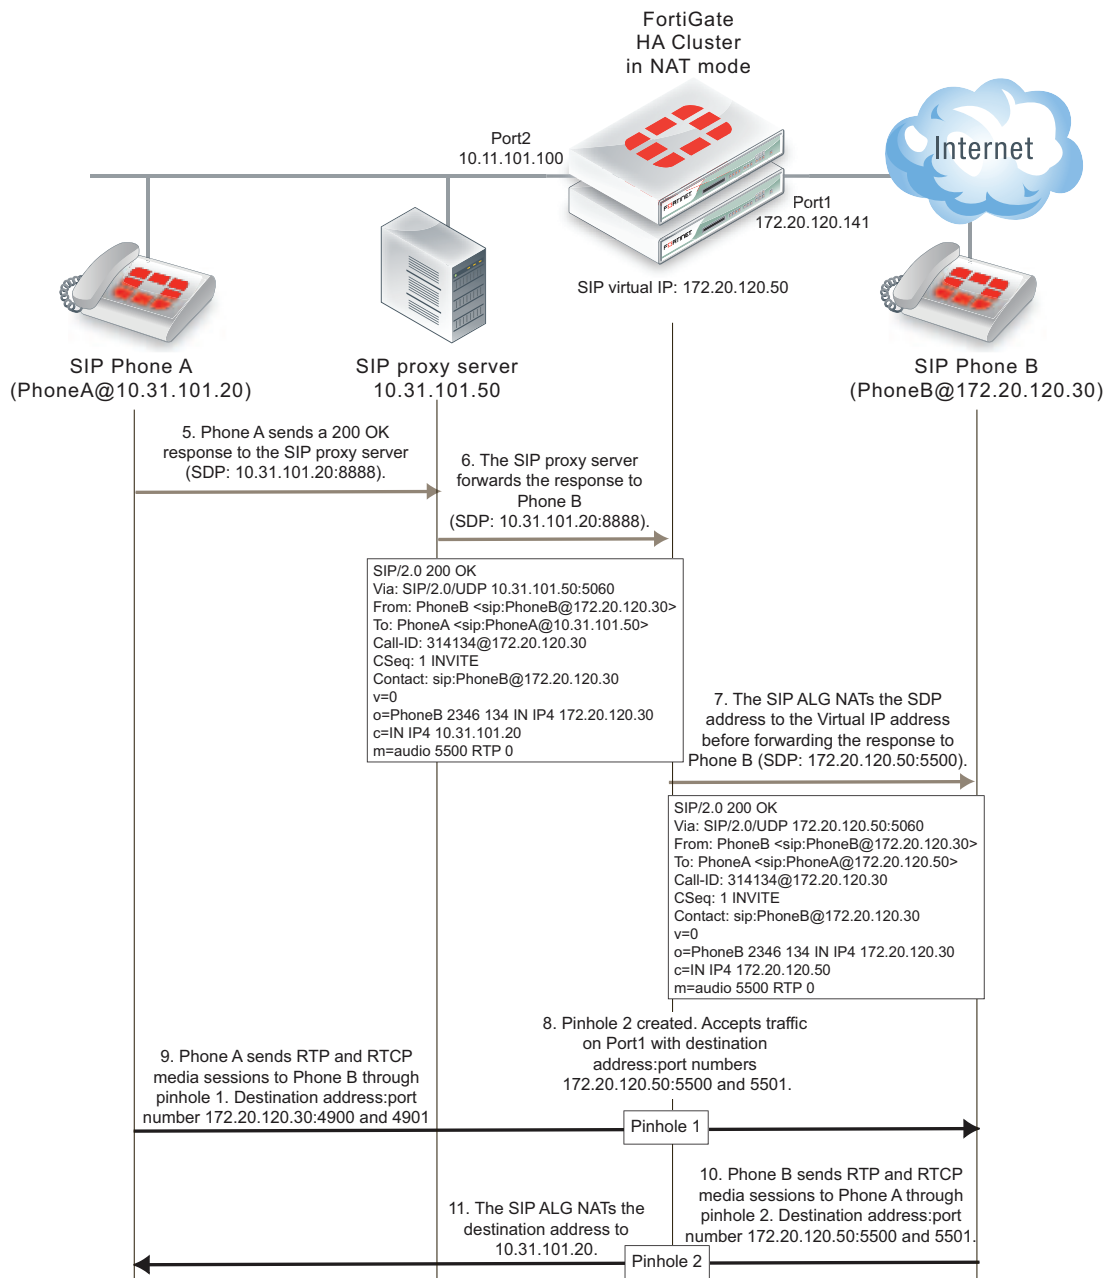
To simplify the diagrams, some SIP messages are not included (for example, the Ringing and ACK response messages) and some SIP header lines and SDP profile lines have been removed from the SIP messages.

The SIP ALG also translates the destination addresses in the SIP message from the virtual IP address (172.20.120.50) to the SIP proxy server address (10.31.101.50). For this configuration to work, the SIP proxy server must be able to change the destination addresses for Phone A in the SIP message from the address of the SIP proxy server to the actual address of Phone A.

The SIP ALG also opens a pinhole on the Port2 interface that accepts media sessions from the private network to SIP Phone B using ports 4900 and 4901.

Phone A sends a 200 OK response back to the SIP proxy server. The SIP proxy server forwards the response to Phone B. The FortiGate accepts the 100 OK response. The SIP ALG translates the Phone A addresses back to the SIP proxy server virtual IP address before forwarding the response back to Phone B. The SIP ALG also opens a pinhole using the SIP proxy server virtual IP which is the address in the o= line of the SDP profile and the port number in the m= line of the SDP code.

SIP destination NAT scenario part 2: 200 OK returned to Phone B and media streams established




The media stream from Phone A is accepted by pinhole one and forwarded to Phone B. The source address of this media stream is changed to the SIP proxy server virtual IP address. The media stream from Phone B is accepted by pinhole 2 and forwarded to Phone B. The destination address of this media stream is changed to the IP address of Phone A.

SIP NAT configuration example: source address translation (source NAT)

This configuration example shows how to configure the FortiGate to support the source address translation scenario shown below. The FortiGate requires two security policies that accept SIP packets. One to allow SIP Phone A to start a session with SIP Phone B and one to allow SIP Phone B to start a session with SIP Phone A. Both of these policies must include source NAT. In this example the networks are not hidden from each other so destination NAT is not required.

SIP source NAT configuration



General configuration steps

The following general configuration steps are required for this SIP configuration. This example uses the default VoIP profile. The example also includes security policies that specifically allow SIP sessions using UDP port 5060 from Phone A to Phone B and from Phone B to Phone A. In most cases you would have more than two phones so would use more general security policies. Also, you can set the firewall service to ANY to allow traffic other than SIP on UDP port 5060.

1. Add firewall addresses for Phone A and Phone B.
2. Add a security policy that accepts SIP sessions initiated by Phone A and includes the default VoIP profile.
3. Add a security policy that accepts SIP sessions initiated by Phone B and includes the default VoIP profile.

Configuration steps - GUI

To add firewall addresses for the SIP phones

1. Go to **Policy & Objects > Addresses**.
2. Add the following addresses for Phone A and Phone B:

Category	Address
Name	Phone_A
Type	IP/Netmask
Subnet / IP Range	10.31.101.20/255.255.255.255
Interface	Internal

Category	Address
Name	Phone_B
Type	IP/Netmask
Subnet / IP Range	172.20.120.30/255.255.255.255
Interface	wan1

To add security policies to apply the SIP ALG to SIP sessions

1. Go to **Policy & Objects > Policy > IPv4**.
2. Add a security policy to allow Phone A to send SIP request messages to Phone B:

Incoming Interface	internal
Outgoing Interface	wan1
Source	Phone_A
Destination Address	Phone_B
Schedule	always
Service	SIP
Action	ACCEPT

3. Turn on **NAT** and select **Use Outgoing Interface Address**.
4. Turn on **VoIP** and select the **default** VoIP profile.
5. Select **OK**.
6. Add a security policy to allow Phone B to send SIP request messages to Phone A:

Incoming Interface	wan1
Outgoing Interface	internal
Source	Phone_B
Destination Address	Phone_A
Schedule	always
Service	SIP
Action	ACCEPT

7. Turn on **NAT** and select **Use Outgoing Interface Address**.
8. Turn on **VoIP** and select the **default** VoIP profile.
9. Select **OK**.

Configuration steps - CLI

To add firewall addresses for Phone A and Phone B and security policies to apply the SIP ALG to SIP sessions

1. Enter the following command to add firewall addresses for Phone A and Phone B.

```
config firewall address
  edit Phone_A
    set associated interface internal
    set type ipmask
    set subnet 10.31.101.20 255.255.255.255
  next
  edit Phone_B
    set associated interface wan1
    set type ipmask
    set subnet 172.20.120.30 255.255.255.255
end
```

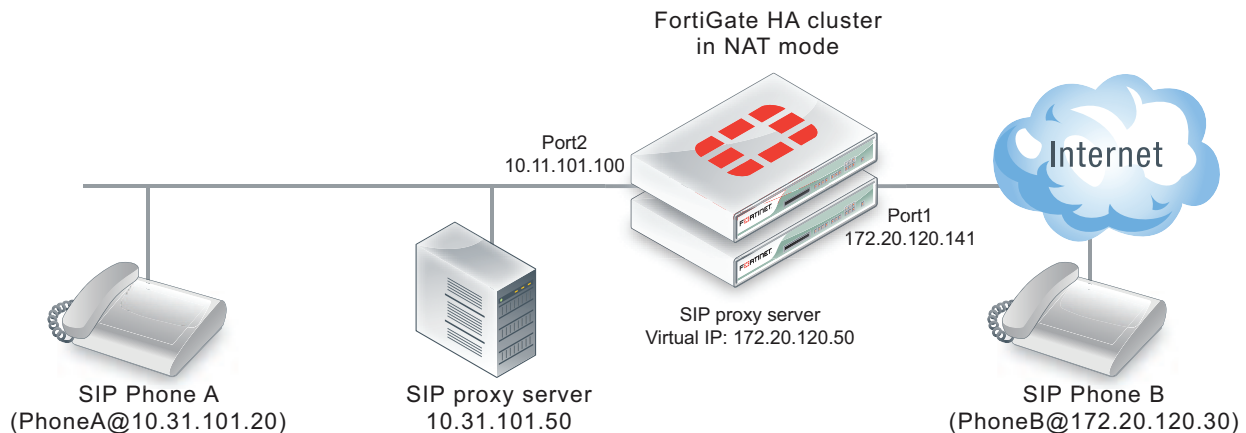
2. Enter the following command to add security policies to allow Phone A to send SIP request messages to Phone B and Phone B to send SIP request messages to Phone A.

```
config firewall policy
  edit 0
    set srcintf internal
    set dstintf wan1
    set srcaddr Phone_A
    set dstaddr Phone_B
    set action accept
    set schedule always
    set service SIP
    set nat enable
    set utm-status enable
    set voip-profile default
  next
  edit 0
    set srcintf wan1
    set dstintf internal
    set srcaddr Phone_B
    set dstaddr Phone_A
    set action accept
    set schedule always
    set service SIP
    set nat enable
    set utm-status enable
    set voip-profile default
end
```

SIP NAT configuration example: destination address translation (destination NAT)

This configuration example shows how to configure the FortiGate to support the destination address translation scenario shown in the figure below. The FortiGate requires two SIP security policies:

- A destination NAT security policy that allows SIP messages to be sent from the Internet to the private network. This policy must include destination NAT because the addresses on the private network are not routable on the Internet.
- A source NAT security policy that allows SIP messages to be sent from the private network to the Internet.

SIP destination NAT scenario part two: 200 OK returned to Phone B and media streams established**General configuration steps**

The following general configuration steps are required for this destination NAT SIP configuration. This example uses the default VoIP profile.

1. Add the SIP proxy server firewall virtual IP.
2. Add a firewall address for the SIP proxy server on the private network.
3. Add a destination NAT security policy that accepts SIP sessions from the Internet destined for the SIP proxy server virtual IP and translates the destination address to the IP address of the SIP proxy server on the private network.
4. Add a security policy that accepts SIP sessions initiated by the SIP proxy server and destined for the Internet.

Configuration steps - GUI**To add the SIP proxy server firewall virtual IP**

1. Go to **Policy & Objects > Virtual IPs**.
2. Add the following SIP proxy server virtual IP.

VIP Type	IPv4
Name	SIP_Proxy_VIP
Interface	port1
Type	Static NAT
External IP Address/Range	172.20.120.50
Mapped IP Address/Range	10.31.101.50

To add a firewall address for the SIP proxy server

1. Go to **Policy & Objects > Addresses**.
2. Add the following for the SIP proxy server:

Address Name	SIP_Proxy_Server
Type	Subnet
Subnet/IP Range	10.31.101.50/255.255.255.255
Interface	port2

To add the security policies

1. Go to **Policy & Objects > IPv4 Policy**.
2. Add a destination NAT security policy that includes the SIP proxy server virtual IP that allows Phone B (and other SIP phones on the Internet) to send SIP request messages to the SIP proxy server.

Incoming Interface	port1
Outgoing Interface	port2
Source	all
Destination Address	SIP_Proxy_VIP
Schedule	always
Service	SIP
Action	ACCEPT

3. Turn on **NAT** and select **Use Outgoing Interface Address**.
4. Turn on **VoIP** and select the **default** VoIP profile.
5. Select **OK**.
6. Add a source NAT security policy to allow the SIP proxy server to send SIP request messages to Phone B and the Internet:

Incoming Interface	port2
Destination Address	all
Source	SIP_Proxy_Server
Schedule	always
Service	SIP
Action	ACCEPT

7. Turn on **NAT** and select **Use OutgoingInterface Address**.
8. Turn on **VoIP** and select the **default** VoIP profile.
9. Select **OK**.

Configuration steps - CLI

To add the SIP proxy server firewall virtual IP and firewall address

1. Enter the following command to add the SIP proxy server firewall virtual IP.

```
config firewall vip
edit SIP_Proxy_VIP
set type static-nat
set extip 172.20.120.50
set mappedip 10.31.101.50
set extintf port1
end
```

2. Enter the following command to add the SIP proxy server firewall address.

```
config firewall address
edit SIP_Proxy_Server
set associated interface port2
set type ipmask
set subnet 10.31.101.50 255.255.255.255
end
```

To add security policies

1. Enter the following command to add a destination NAT security policy that includes the SIP proxy server virtual IP that allows Phone B (and other SIP phones on the Internet) to send SIP request messages to the SIP proxy server.

```
config firewall policy
edit 0
set srcintf port1
set dstintf port2
set srcaddr all
set dstaddr SIP_Proxy_VIP
set action accept
set schedule always
set service SIP
set nat enable
set utm-status enable
set voip-profile default
end
```

2. Enter the following command to add a source NAT security policy to allow the SIP proxy server to send SIP request messages to Phone B and the Internet:


```
config firewall policy
edit 0
set srcintf port2
set dstintf port1
set srcaddr SIP_Proxy_Server
set dstaddr all
set action accept
set schedule always
set service SIP
set nat enable
set utm-status enable
set voip-profile default
end
```

SIP and RTP source NAT

In the source NAT scenario shown below, a SIP phone connects to the Internet through a FortiGate with an IP address configured using PPPoE. The SIP ALG translates all private IPs in the SIP contact header into public IPs.

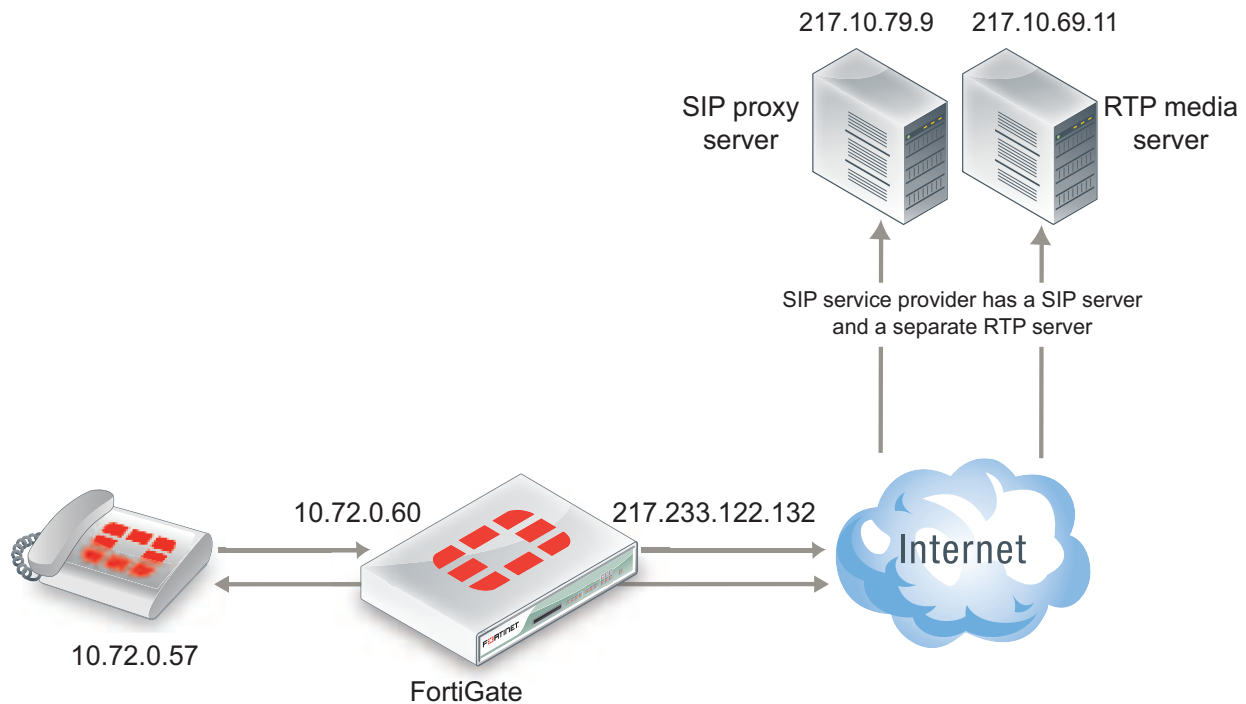
You need to configure an internal to external SIP security policy with NAT selected, and include a VoIP profile with SIP enabled.

SIP source NAT



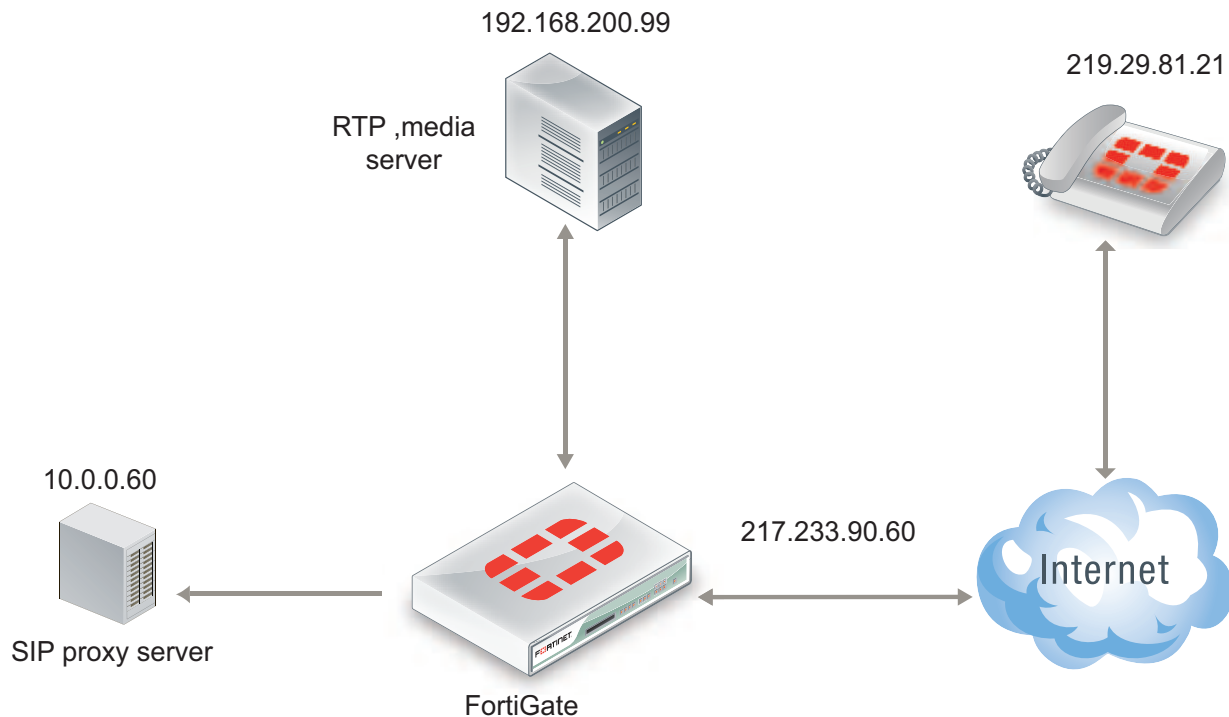
SIP and RTP destination NAT

In the following destination NAT scenario, a SIP phone can connect through the FortiGate to a private IP address using a firewall virtual IP (VIP). The SIP ALG translates the SIP contact header to the IP of the real SIP proxy server located on the Internet.

SIP destination NAT

In the scenario, shown above, the SIP phone connects to a VIP (10.72.0.60). The SIP ALG translates the SIP contact header to 217.10.79.9, opens RTP pinholes, and manages NAT.

The FortiGate also supports a variation of this scenario where the RTP media server's IP address is hidden on a private network or DMZ.

SIP destination NAT-RTP media server hidden

In the scenario shown above, a SIP phone connects to the Internet. The VoIP service provider only publishes a single public IP. The FortiGate is configured with a firewall VIP. The SIP phone connects to the FortiGate (217.233.90.60) and using the VIP the FortiGate translates the SIP contact header to the SIP proxy server IP address (10.0.0.60). The SIP proxy server changes the SIP/SDP connection information (which tells the SIP phone which RTP media server IP it should contact) also to 217.233.90.60.

Source NAT with an IP pool

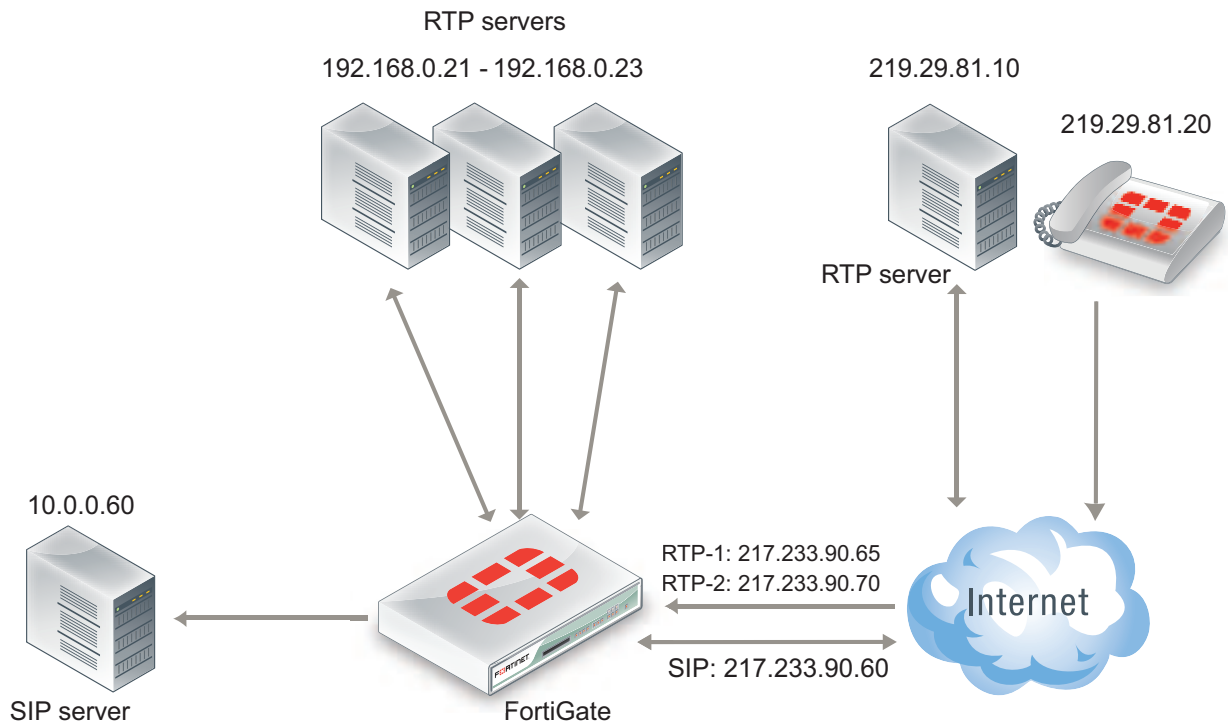
You can choose **NAT** with the **Dynamic IP Pool** option when configuring a security policy if the source IP of the SIP packets is different from the interface IP. The FortiGate ALG interprets this configuration and translates the SIP header accordingly.

This configuration also applies to destination NAT.

Different source and destination NAT for SIP and RTP

This is a more complex scenario that a SIP service provider may use. It can also be deployed in large-scale SIP environments where RTP has to be processed by the FortiGate and the RTP server IP has to be translated differently than the SIP server IP.

Different source and destination NAT for SIP and RTP



In this scenario, shown above, assume there is a SIP server and a separate media gateway. The SIP server is configured so that the SIP phone (219.29.81.20) will connect to 217.233.90.60. The media gateway (RTP server: 219.29.81.10) will connect to 217.233.90.65.

What happens is as follows:

1. The SIP phone connects to the SIP VIP. The FortiGate ALG translates the SIP contact header to the SIP server: 219.29.81.20 > 217.233.90.60 (> 10.0.0.60).
2. The SIP server carries out RTP to 217.233.90.65.
3. The FortiGate ALG opens pinholes, assuming that it knows the ports to be opened.
4. RTP is sent to the RTP-VIP (217.233.90.65.) The FortiGate ALG translates the SIP contact header to 192.168.0.21.

NAT with IP address conservation

In a source or destination NAT security policy that accepts SIP sessions, you can configure the SIP ALG or the SIP session helper to preserve the original source IP address of the SIP message in the `i=` line of the SDP profile. NAT with IP address conservation (also called SIP NAT tracing) changes the contents of SIP messages by adding the source IP address of the originator of the message into the SDP `i=` line of the SIP message. The SDP `i=` line is used for free-form text. However, if your SIP server can retrieve information from the SDP `i=` line, it can be useful for keeping a record of the source IP address of the originator of a SIP message when operating in a NAT environment. You can use this feature for billing purposes by extracting the IP address of the originator of the message.

Configuring SIP IP address conservation for the SIP ALG

You can use the following command to enable or disable SIP IP address conservation in a VoIP profile for the SIP ALG. SIP IP address conservation is enabled by default in a VoIP profile.

```
config voip profile
  edit VoIP_Pro_1
    config sip
      set nat-trace disable
    end
  end
```

If the SIP message does not include an `i=` line and if the original source IP address of the traffic (before NAT) was 10.31.101.20 then the FortiGate would add the following `i=` line.

```
i=(o=IN IP4 10.31.101.20)
```

You can also use the `preserve-override` option to configure the SIP ALG to either add the original `o=` line to the end of the `i=` line or replace the `i=` line in the original message with a new `i=` line in the same form as above for adding a new `i=` line.

By default, `preserve-override` is disabled and the SIP ALG adds the original `o=` line to the end of the original `i=` line. Use the following command to configure the SIP ALG to replace the original `i=` line:

```
config voip profile
  edit VoIP_Pro_1
    config sip
      set preserve-override enable
    end
  end
```

Configuring SIP IP address conservation for the SIP session helper

You can use the following command to enable or disable SIP IP address conservation for the SIP session helper. IP address conservation is enabled by default for the SIP session helper.

```
config system settings
  set sip-nat-trace disable
end
```

If the SIP message does not include an `i=` line and if the original source IP address of the traffic (before NAT) was 10.31.101.20 then the FortiGate would add the following `i=` line.

```
i=(o=IN IP4 10.31.101.20)
```

Controlling how the SIP ALG NATs SIP contact header line addresses

You can enable `contact-fixup` so that the SIP ALG performs normal SIP NAT translation to SIP contact headers as SIP messages pass through the FortiGate.

Disable `contact-fixup` if you do not want the SIP ALG to perform normal NAT translation of the SIP contact header if a Record-Route header is also available. If `contact-fixup` is disabled, the FortiGate ALG does the following with contact headers:

- For Contact in Requests, if a Record-Route header is present and the request comes from the external network, the SIP Contact header is not translated.
- For Contact in Responses, if a Record-Route header is present and the response comes from the external network, the SIP Contact header is not translated.

If `contact-fixup` is disabled, the SIP ALG must be able to identify the external network. To identify the external network, you must use the `config system interface` command to set the `external` keyword to `enable` for the interface that is connected to the external network.

Enter the following command to perform normal NAT translation of the SIP contact header:

```
config voip profile
  edit VoIP_Pro_1
    config sip
      set contact-fixup enable
    end
  end
```

Controlling NAT for addresses in SDP lines

You can use the `no-sdp-fixup` option to control whether the FortiGate performs NAT on addresses in SDP lines in the SIP message body.

The `no-sdp-fixup` option is disabled by default and the FortiGate performs NAT on addresses in SDP lines. Enable this option if you don't want the FortiGate to perform NAT on the addresses in SDP lines.


```
config voip profile
  edit VoIP_Pro_1
    config sip
      set no-sdp-fixup enable
    end
  end
```

Translating SIP session destination ports

Using port forwarding virtual IPs you can change the destination port of SIP sessions as they pass through the FortiGate.

Translating SIP sessions to a different destination port

To configure translating SIP sessions to a different destination port you must add a static NAT virtual IP that translates the SIP destination port to another port destination. In the example the destination port is translated from 5060 to 50601. This configuration can be used if SIP sessions use different destination ports on different networks.

Example translating SIP sessions to a different destination port**To translate SIP sessions to a different destination port**

1. Add the static NAT virtual IP.

This virtual IP forwards traffic received at the port1 interface for IP address 172.20.120.20 and destination port 5060 to the SIP server at IP address 192.168.10.20 with destination port 5061.

```
config firewall vip
  edit "sip_port_trans_vip"
    set type static-nat
    set portforward enable
    set protocol tcp
    set extip 172.20.120.20
    set extport 5060
    set extintf "port1"
    set mappedip 192.168.10.20
    set mappedport 50601
    set comment "Translate SIP destination port"
  end
```

2. Add a security policy that includes the virtual IP and the default VoIP profile.

```
config firewall policy
  edit 1
    set srcintf "port1"
    set dstintf "port2"
    set srcaddr "all"
```

```


set dstaddr "sip_port_trans_vip"
set action accept
set schedule "always"
set service "ANY"
set utm-status enable
set profile-protocol-options default
set comments "Translate SIP destination port"
end

```

Translating SIP sessions to multiple destination ports

You can use a load balance virtual IP to translate SIP session destination ports to a range of destination ports. In this example the destination port is translated from 5060 to the range 50601 to 50603. This configuration can be used if your SIP server is configured to receive SIP traffic on multiple ports.

Example translating SIP traffic to multiple destination ports



To translated SIP sessions to multiple destination ports

1. Add the load balance virtual IP.

This virtual IP forwards traffic received at the port1 interface for IP address 172.20.120.20 and destination port 5060 to the SIP server at IP address 192.168.10.20 with destination port 5061.

```
config firewall vip
```

```

edit "sip_port_ldbl_vip"
    set type load-balance
    set portforward enable
    set protocol tcp
    set extip 172.20.120.20
    set extport 5060
    set extintf "port1"
    set mappedip 192.168.10.20
    set mappedport 50601-50603
    set comment "Translate SIP destination port range"
end

```

2. Add a security policy that includes the virtual IP and VoIP profile.

```

config firewall policy
    edit 1
        set srcintf "port1"
        set dstintf "port2"
        set srcaddr "all"
        set dstaddr "sip_port_ldbl_vip"
        set action accept
        set schedule "always"
        set service "ANY"
        set utm-status enable
        set voip-profile default
        set comments "Translate SIP destination port"
    end

```

Adding the original IP address and port to the SIP message header after NAT

In some cases your SIP configuration may require that the original IP address and port from the SIP contact request is kept after NAT. For example, the original SIP contact request could include the following:

```
Contact: <sip:0150302438@172.20.120.110:5060>;
```

After the packet goes through the FortiGate and NAT is performed, the contact request could normally look like the following (the IP address translated to a different IP address and the port to a different port):

```
Contact: <sip:0150302438@10.10.10.21:33608>;
```

You can enable `register-contact-trace` in a VoIP profile to have the SIP ALG add the original IP address and port in the following format:

```
Contact: <sip:0150302438@<nated-ip>:<nated-port>;o=<original-ip>: <original-port>>;
```

So the contact line after NAT could look like the following:

```
Contact: <sip:0150302438@10.10.10.21:33608;o=172.20.120.110:5060>;
```

Enter the following command to enable keeping the original IP address and port:

```

config voip profile
    edit Profile_name
        config sip
            set register-contract-trace enable
        end
    end

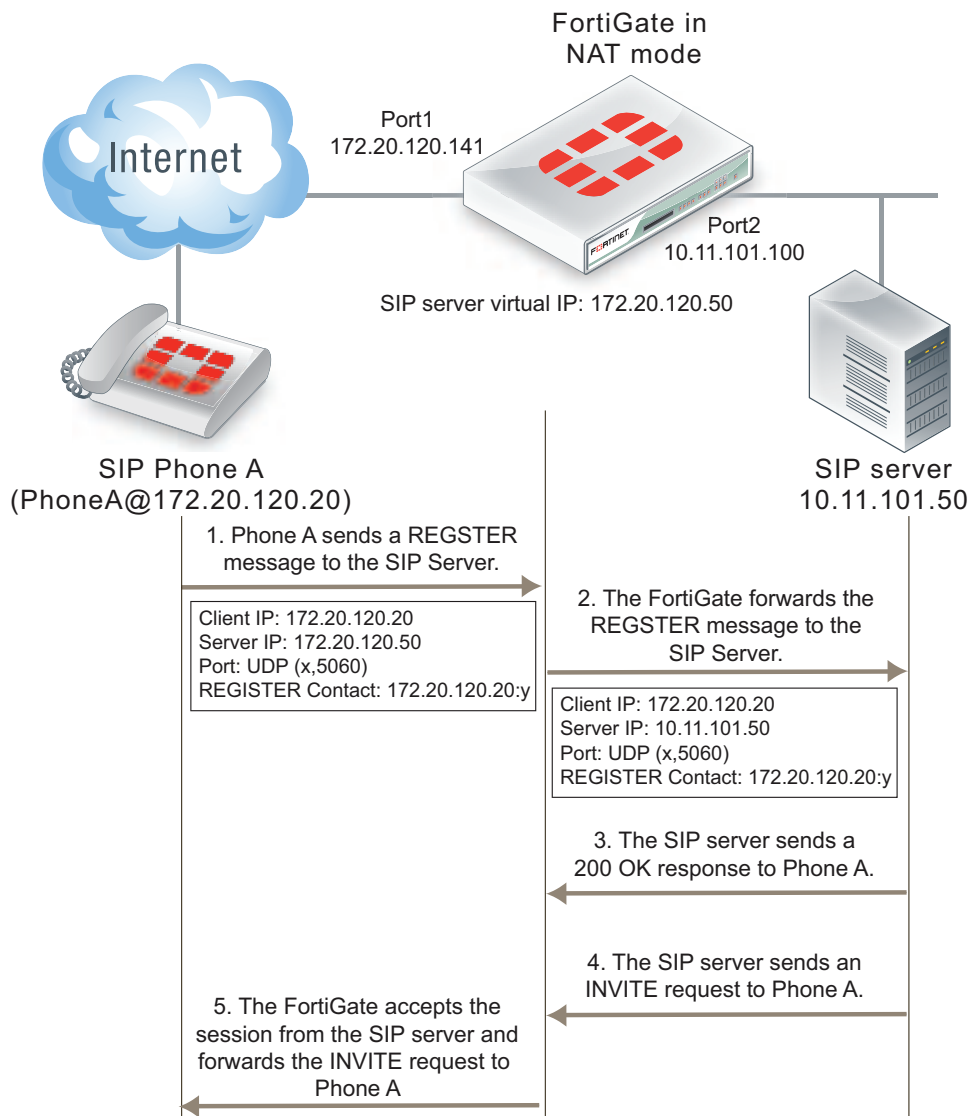
```

Enhancing SIP pinhole security

You can use the `strict-register` option in a SIP VoIP profile to open smaller pinholes. This option is enabled by default on the default VoIP profiles and in all new VoIP profiles that you create.

As shown below, when FortiGate is protecting a SIP server on a private network, the FortiGate does not have to open a pinhole for the SIP server to send INVITE requests to a SIP Phone on the Internet after the SIP Phone has registered with the server.

FortiGate protecting a SIP server on a private network



In the example, a client (SIP Phone A) sends a REGISTER request to the SIP server with the following information:

```
Client IP: 10.31.101.20
Server IP: 10.21.101.50
```

```
Port: UDP (x,5060)
REGISTER Contact: 10.31.101.20:y
```

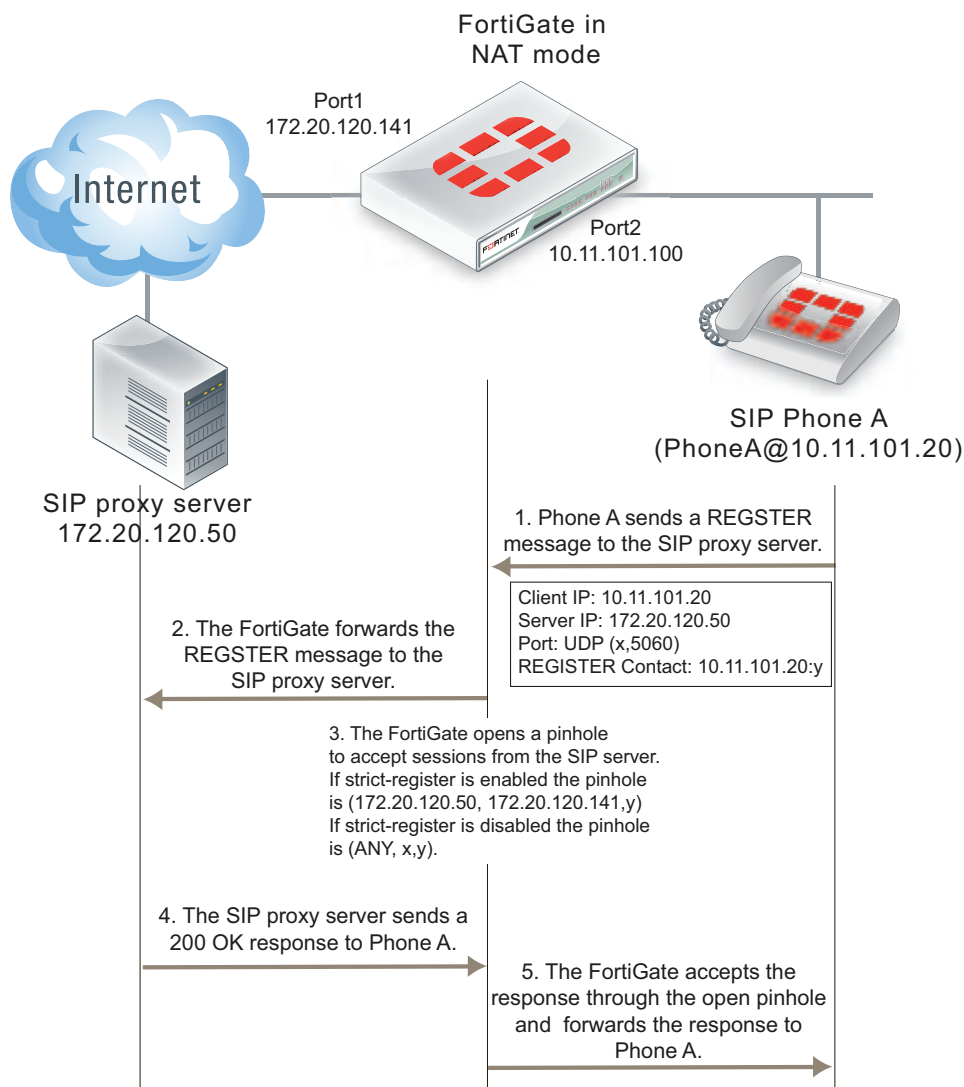
Where *x* and *y* are ports chosen by Phone A.

As soon as the server sends the 200 OK reply it can forward INVITE requests from other SIP phones to SIP Phone A. If the SIP proxy server uses the information in the REGISTER message received from SIP Phone A the INVITE messages sent to Phone A will only get through the FortiGate if a policy has been added to allow the server to send traffic from the private network to the Internet. Or the SIP ALG must open a pinhole to allow traffic from the server to the Internet. In most cases the FortiGate is protecting the SIP server so there is no reason not to add a security policy to allow the SIP server to send outbound traffic to the Internet.

In a typical SOHO scenario, shown below, SIP Phone A is being protected from the Internet by a FortiGate. In most cases the FortiGate would not allow incoming traffic from the Internet to reach the private network. So the only way that an INVITE request from the SIP server can reach SIP Phone A is if the SIP ALG creates an incoming pinhole. All pinholes have three attributes:

```
(source address, destination address, destination port)
```

SOHO configuration, FortiGate protecting a network with SIP phones



The more specific a pinhole is the more secure it is because it accept less traffic. In this situation, the pinhole would be more secure if it only accepted traffic from the SIP server. This is what happens if `strict-register` is enabled in the VoIP profile that accepts the REGISTER request from Phone A.

(SIP server IP address, client IP address, destination port)

If `strict-register` is disabled (the default configuration) the pinhole is set up with the following attributes

(ANY IP address, client IP address, destination port)

This pinhole allows connections through the FortiGate from ANY source address which is a much bigger and less secure pinhole. In most similar network configurations you should enable `strict-register` to improve pinhole security.

Enabling `strict-register` can cause problems when the SIP registrar and SIP proxy server are separate entities with separate IP addresses.

Enter the following command to enable `strict-register` in a VoIP profile.

```
config voip profile
  edit Profile_name
    config sip
      set strict-register enable
    end
```

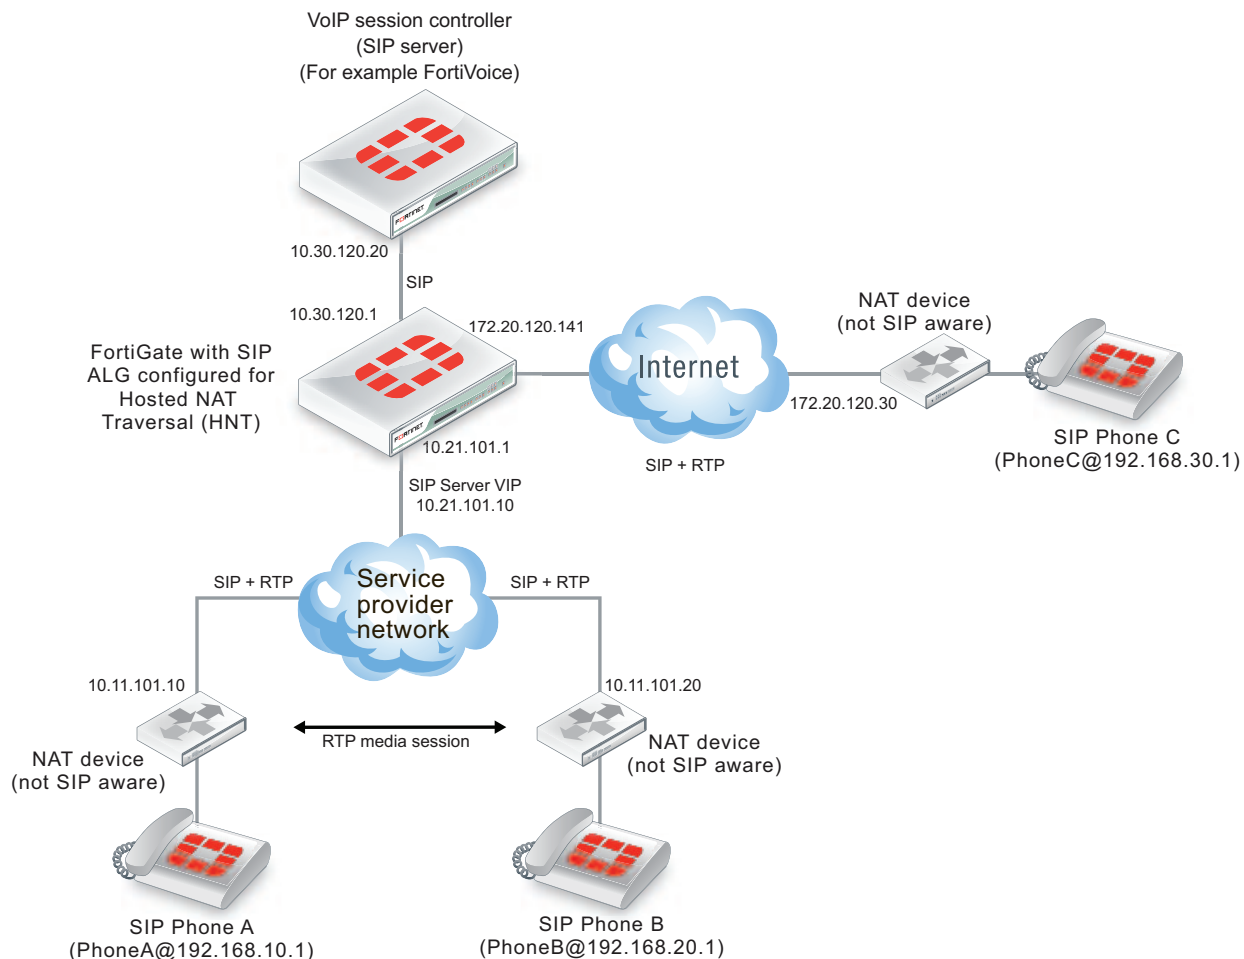
Hosted NAT traversal

With the increase in the use of VoIP and other media traffic over the Internet, service provider network administrators must defend their networks from threats while allowing voice and multimedia traffic to flow transparently between users and servers and among users. A common scenario could involve providing SIP VoIP services for customers with SIP phones installed behind NAT devices that are not SIP aware. NAT devices that are not SIP aware cannot translate IP addresses in SIP headers and SDP lines in SIP packets but can and do perform source NAT on the source or addresses of the packets. In this scenario the user's SIP phones would communicate with a SIP proxy server to set up calls between SIP phones. Once the calls are set up RTP packets would be communicated directly between the phones through each user's NAT device.

The problem with this configuration is that the SIP headers and SDP lines in the SIP packets sent from the phones and received by the SIP proxy server would contain the private network addresses of the VoIP phones that would not be routable on the service provider network or on the Internet. One solution could be to for each customer to install and configure SIP aware NAT devices. If this is not possible, another solution requires implement hosted NAT traversal.

In a hosted NAT traversal (HNT) configuration, a FortiGate is installed between the NAT device and the SIP proxy server and configured with a VoIP profile that enables SIP hosted NAT traversal. Security policies that include the VoIP profile also support destination NAT using a firewall virtual IP. When the SIP phones connect to the SIP server IP address the security policy accepts the SIP packets, the virtual IP translates the destination addresses of the packets to the SIP server IP address, and the SIP ALG NAT traversal configuration translates the source IP addresses on the SIP headers and SDP lines to the source address of the SIP packets (which would be the external IP address of the NAT devices). The SIP server then sees the SIP phone IP address as the external IP address of the NAT device. As a result SIP and RTP media sessions are established using the external IP addresses of the NAT devices instead of the actual IP addresses of the SIP phones.

FortiGate SIP Hosted NAT Traversal configuration



Configuration example: Hosted NAT traversal for calls between SIP Phone A and SIP Phone B

The following address translation takes place to allow a SIP call from SIP Phone A to SIP Phone B in the above diagram.

1. SIP Phone A sends a SIP Invite message to the SIP server. Packet source IP address: 192.168.10.1, destination IP address: 10.21.101.10.
2. The SIP packets are received by the NAT device which translates the source address of the SIP packets from 192.168.10.1 to 10.11.101.20.
3. The SIP packets are received by the FortiGate which translates the packet destination IP address to 10.30.120.20. The SIP ALG also translates the IP address of the SIP phone in the SIP header and SDP lines from 192.168.10.1 to 10.11.101.20.
4. The SIP server accepts the Invite message and forwards it to SIP Phone B at IP address 10.11.101.20. The SIP server has this address for SIP Phone B because SIP packets from SIP Phone B have also been translated using the hosted NAT traversal configuration of the SIP ALG.

5. When the SIP call is established, the RTP session is between 10.11.101.10 and 10.11.101.20 and does not pass through the FortiGate. The NAT devices translated the destination address of the RTP packets to the private IP addresses of the SIP phones.

General configuration steps

The following general configuration steps are required for this destination NAT SIP configuration. This example uses the default VoIP profile.

1. Add a VoIP profile that enables hosted NAT translation.
2. Add a SIP proxy server firewall virtual IP.
3. Add a firewall address for the SIP proxy server on the private network.
4. Add a destination NAT security policy that accepts SIP sessions from the Internet destined for the SIP proxy server virtual IP and translates the destination address to the IP address of the SIP proxy server on the private network.
5. Add a security policy that accepts SIP sessions initiated by the SIP proxy server and destined for the Internet.

Configuration steps - GUI

To add the SIP proxy server firewall virtual IP

1. Go to **Policy & Objects > Virtual IPs**.
2. Add the SIP proxy server virtual IP.

Name	SIP_Proxy_VIP
External Interface	port1
Type	Static NAT
External IP Address/Range	172.20.120.50
Mapped IP Address/Range	10.31.101.50

To add a firewall address for the SIP proxy server

1. Go to **Policy & Objects > Addresses**.
2. Add the following for the SIP proxy server:

Category	Address
Name	SIP_Proxy_Server
Type	Subnet
Subnet / IP Range	10.31.101.50/255.255.255.255
Interface	port2

To add the security policies

1. Go to **Policy & Objects > IPv4 Policy**.
2. Add a destination NAT security policy that includes the SIP proxy server virtual IP that allows Phone B (and other SIP phones on the Internet) to send SIP request messages to the SIP proxy server.

Incoming Interface	port1
Outgoing Interface	port2
Source	all
Destination Address	SIP_Proxy_VIP
Schedule	always
Service	SIP
Action	ACCEPT

3. Turn on **NAT** and select **Use Outgoing Interface Address**.
4. Turn on **VoIP** and select the **default** VoIP profile.
5. Select **OK**.
6. Add a source NAT security policy to allow the SIP proxy server to send SIP request messages to Phone B and the Internet:

Incoming Interface	port2
Outgoing Interface	port1
Source	SIP_Proxy_Server
Destination Address	all
Schedule	always
Service	SIP
Action	ACCEPT

7. Turn on **NAT** and select **Use Outgoing Interface Address**.
8. Turn on **VoIP** and select the **default** VoIP profile.
9. Select **OK**.

Configuration steps - CLI**To add a VoIP profile that enables hosted NAT translation**

1. Enter the following command to add a VoIP profile named HNT that enables hosted NAT traversal. This command shows how to clone the default VoIP profile and enable hosted NAT traversal.

```
config voip profile
  clone default to HNT
```

```
edit HNT
  config sip
    set hosted-nat-traversal enable
  end
end
```

To add the SIP proxy server firewall virtual IP and firewall address

1. Enter the following command to add the SIP proxy server firewall virtual IP.

```
config firewall vip
  edit SIP_Proxy_VIP
    set type static-nat
    set extip 10.21.101.10
    set mappedip 10.30.120.20
    set extintf port1
  end
```

2. Enter the following command to add the SIP proxy server firewall address.

```
config firewall address
  edit SIP_Proxy_Server
    set associated interface port2
    set type ipmask
    set subnet 10.30.120.20 255.255.255.255
  end
```

To add security policies

1. Enter the following command to add a destination NAT security policy that includes the SIP proxy server virtual IP that allows Phone A to send SIP request messages to the SIP proxy server.

```
config firewall policy
  edit 0
    set srcintf port1
    set dstintf port2
    set srcaddr all
    set dstaddr SIP_Proxy_VIP
    set action accept
    set schedule always
    set service SIP
    set nat enable
    set utm-status enable
    set profile-protocol-options default
    set voip-profile HNT
  end
```

2. Enter the following command to add a source NAT security policy to allow the SIP proxy server to send SIP request messages to Phone B:

```
config firewall policy
  edit 0
    set srcintf port2
    set dstintf port1
    set srcaddr SIP_Proxy_Server
    set dstaddr all
    set action accept
    set schedule always
    set service SIP
    set nat enable
```

```
set utm-status enable
set profile-protocol-options default
set voip-profile default
end
```

Hosted NAT traversal for calls between SIP Phone A and SIP Phone C

The following address translation takes place to allow a SIP call from SIP Phone A to SIP Phone C in the previous diagram.

1. SIP Phone A sends a SIP Invite message to the SIP server. Packet source IP address: 192.168.10.1 and destination IP address: 10.21.101.10.
2. The SIP packets are received by the NAT device which translates the source address of the SIP packets from 192.168.10.1 to 10.11.101.10.
3. The SIP packets are received by the FortiGate which translates the packet destination IP address to 10.30.120.20. The SIP ALG also translates the IP address of the SIP phone in the SIP header and SDP lines from 192.168.10.1 to 10.11.101.20.
4. The SIP server accepts the Invite message and forwards it to SIP Phone C at IP address 172.20.120.30. The SIP server has this address for SIP Phone C because SIP packets from SIP Phone C have also been translated using the hosted NAT traversal configuration of the SIP ALG.
5. When the SIP call is established, the RTP session is between 10.11.101.10 and 172.20.120.30. The packets pass through the FortiGate which performs NAT as required.

Restricting the RTP source IP

Use the following command in a VoIP profile to restrict the RTP source IP to be the same as the SIP source IP when hosted NAT traversal is enabled.

```
config voip profile
  edit VoIP_HNT
    config sip
      set hosted-nat-traversal enable
      set hnt-restrict-source-ip enable
    end
  end
```


SIP over IPv6

FortiGates operating in NAT/Route and in transparent mode support SIP over IPv6. The SIP ALG can process SIP messages that use IPv6 addresses in the headers, bodies, and in the transport stack. The SIP ALG cannot modify the IPv6 addresses in the SIP headers so FortiGates cannot perform SIP or RTP NAT over IPv6 and also cannot translate between IPv6 and IPv4 addresses.

In the scenario shown in the figure below, a SIP phone connects to the Internet through a FortiGate operating. The phone and the SIP and RTP servers all have IPv6 addresses.

The FortiGate has IPv6 security policies that accept SIP sessions. The SIP ALG understands IPv6 addresses and can forward IPv6 sessions to their destinations. Using SIP application control features the SIP ALG can also apply rate limiting and other settings to SIP sessions.

SIP support for IPv6

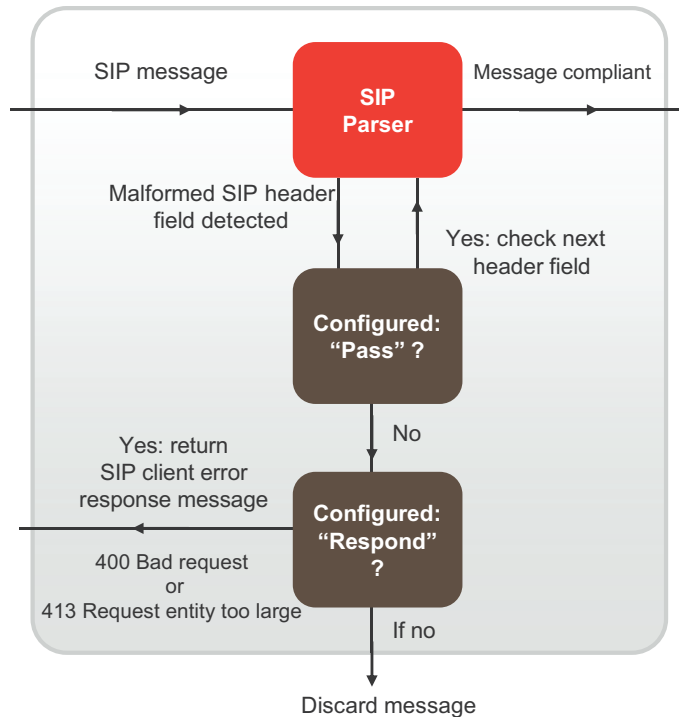


To enable SIP support for IPv6 add an IPv6 security policy that accepts SIP packets and includes a VoIP profile.

Deep SIP message inspection

Deep SIP message syntax inspection (also called Deep SIP header inspection or SIP fuzzing protection) provides protection against malicious SIP messages by applying SIP header and SDP profile syntax checking. SIP Fuzzing attacks can be used by attackers to discover and exploit vulnerabilities of a SIP entity (for example a SIP proxy server). Most often these attacks could crash or compromise the SIP entity.

Deep SIP message inspection



- Checks the SIP request message Request-line
- Checks the following SIP header fields:
 - Allow, Call-id, Contact, Content-length, Content-type, CSeq, Expires, From, Max-Forwards, P-asserted-identity, Rack, Record-Route, Route, Rseq, To, Via
- Checks all SDP profile lines
- Configurable header and body length checks
- Optional logging of message violations

Deep SIP message inspection checks the syntax of each SIP header and SDP profile line to make sure they conform to the syntax defined in the relevant RFC and IETF standard. You can also configure the SIP ALG to inspect for:

- Unknown SIP message types (message types not defined in a SIP RFC) this option is enabled by default and can be disabled. When enabled unknown message types are discarded. Configured using the `block-unknown` option.
- Unknown line types (message line types that are not defined in any SIP or SDP RFC). Configured using the `unknown-header` option.
- Messages that are longer than a configured maximum size. Configured using the `max-body-length` option.
- Messages that contain one or more lines that are longer than a set maximum line length (default 998 characters). Configured using the `max-line-length` option.

Actions taken when a malformed message line is found

When a malformed message line or other error is found the SIP ALG can be configured to discard the message containing the error, pass the message without any other actions, or responding to the message with a 400 Bad Request or 413 Request entity too large client error SIP response message and then discard the message. (For information about client error SIP response messages, see [Client error on page 3108](#).)

If a message line is longer than the configured maximum, the SIP ALG sends the following message:

```
SIP/2.0 413 Request Entity Too Large, <optional_info>
```

If a message line is incorrect or in an unknown message line is found, the SIP ALG sends the following message:

```
SIP/2.0 400 Bad Request, <optional_info>
```

The `<optional_info>` provides more information about why the message was rejected. For example, if the SIP ALG finds a malformed Via header line, the response message may be:

```
SIP/2.0 400 Bad Request, malformed Via header
```

If the SIP ALG finds a malformed message line, and the action for this message line type is discard, the message is discarded with no further checking or responses. If the action is pass, the SIP ALG continues parsing the SIP message for more malformed message lines. If the action is respond, the SIP ALG sends the SIP response message and discards the message containing the malformed line with no further checking or response. If only malformed message line types with action set to pass are found, the SIP ALG extracts as much information as possible from the message (for example for NAT and opening pinholes, and forwards the message to its destination).

If a SIP message containing a malformed line is discarded the SIP ALG will not use the information in the message for call processing. This could result in the call being terminated. If a malformed line in a SIP message includes information required for the SIP call that the SIP ALG cannot interpret (for example, if an IP address required for SIP NAT is corrupted) the SIP ALG may not be able to continue processing the call and it could be terminated. Discarded messages are counted by SIP ALG static message counters.

Logging and statistics

To record a log message each time the SIP ALG finds a malformed header, enable logging SIP violations in a VoIP profile. In all cases, when the SIP ALG finds an error the FortiGate records a malformed header log message that contains information about the error. This happens even if the action is set to pass.

If, because of recording log messages for deep message inspection, the CPU performance is affected by a certain amount, the FortiGate records a critical log message about this event and stops writing log messages for deep SIP message inspection.

The following information is recorded in malformed header messages:

- The type of message line in which the error was found.
- The content of the message line in which the error was found (it will be truncated if it makes the log message too long)
- The column or character number in which the error was found (to make it easier to determine what caused the error)

Deep SIP message inspection best practices

Because of the risks imposed by SIP header attacks or incorrect data being allowed and because selecting drop or respond does not require more CPU overhead than pass you would want to set all tests to drop or respond.

However, in some cases malformed lines may be less of a threat or risk. For example, the SDP `i=` does not usually contain information that is parsed by any SIP device so a malformed `i=` line may not pose a threat.

You can also use the pre-defined VoIP profiles to apply different levels of deep message inspection. The default VoIP profile sets all deep message inspection options to pass and the strict VoIP profile sets all deep message inspection options to discard. From the CLI you can use the `clone` command to copy these pre-defined VoIP profiles and then customize them for your requirements.

Configuring deep SIP message inspection

You configure deep SIP message inspection in a VoIP profile. All deep SIP message inspection options are available only from the CLI.

Enter the following command to configure deep SIP message inspection to discard messages with malformed Request-lines (the first line in a SIP request message):

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set malformed-request-line respond
    end
  end
```



You cannot configure message inspection for the Status-line, which is the first line in a SIP response message.

The following table lists the SIP header lines that the SIP ALG can inspect and the CLI command for configuring the action for each line type. The table also lists the RFC that the header line is defined in.

SIP header lines that the SIP ALG can inspect for syntax errors

SIP Header line	VoIP profile option	RFC
Allow	<code>malformed-header-allow</code>	RFC 3261
Call-ID	<code>malformed-header-call-id</code>	RFC 3261
Contact	<code>malformed-header-contact</code>	RFC 3261
Content-Length	<code>malformed-header-content-length</code>	RFC 3261
Content-Type	<code>malformed-header-content-type</code>	RFC 3261
CSeq	<code>malformed-header-cseq</code>	RFC 3261
Expires	<code>malformed-header-expires</code>	RFC 3261
From	<code>malformed-header-from</code>	RFC 3261

SIP Header line	VoIP profile option	RFC
Max-forwards	malformed-header-max-forwards	RFC 3261
P-Asserted-Identity	malformed-header-p-asserted-identity	RFC 3325
RAck	malformed-header-rack	RFC 3262
Record-Route	malformed-header-record-route	RFC 3261
Route	malformed-header-route	RFC 3261
RSeq	malformed-header-rseq	RFC 3262
To	malformed-header-to	RFC 3261
Via	malformed-header-via	RFC 3261

The table below lists the SDP profile lines that the SIP ALG inspects and the CLI command for configuring the action for each line type. SDP profile lines are defined by RFC 4566 and RFC 2327.

SDP profile lines that the SIP ALG can inspect for syntax errors

Attribute	VoIP profile option
a=	malformed-header-sdb-a
b=	malformed-header-sdp-b
c=	malformed-header-sdp-c
i=	malformed-header-sdp-i
k=	malformed-header-sdp-k
m=	malformed-header-sdp-m
o=	malformed-header-sdp-o
r=	malformed-header-sdp-r
s=	malformed-header-sdp-s
t=	malformed-header-sdp-t
v=	malformed-header-sdp-v
z=	malformed-header-sdp-z

Discarding SIP messages with some malformed header and body lines

Enter the following command to configure deep SIP message inspection to discard SIP messages with a malformed Via line, a malformed route line or a malformed m= line but to pass messages with a malformed i= line or a malformed Max-Forwards line

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set malformed-header-via discard
      set malformed-header-route discard
      set malformed-header-sdp-m discard
      set malformed-header-sdp-i pass
      set malformed-header-max-forwards pass
    end
  end
```

Discarding SIP messages with an unknown SIP message type

Enter the following command to discard SIP messages with an unknown SIP message line type as defined in all current SIP RFCs:

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set unknown-header discard
    end
  end
```

Discarding SIP messages that exceed a message size

Enter the following command to set the maximum size of a SIP message to 200 bytes. Messages longer than 200 bytes are discarded.

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set max-body-length 200
    end
  end
```

The `max-body-length` option checks the value in the SIP Content-Length header line to determine body length. The Content-Length can be larger than the actual size of a SIP message if the SIP message content is split over more than one packet. SIP message sizes vary widely. The size of a SIP message can also change with the addition of Via and Record-Route headers as the message is transmitted between users and SIP servers.

Discarding SIP messages with lines longer than 500 characters

Enter the following command to set the length of a SIP message line to 500 characters and to block messages that include lines with 500 or more characters:

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set max-line-length 500
      set block-long-lines enable
    end
  end
```

Blocking SIP request messages

You may want to block different types of SIP requests:

- to prevent SIP attacks using these messages.
- If your SIP server cannot process some SIP messages because of a temporary issue (for example a bug that crashes or compromises the server when it receives a message of a certain type).
- Your SIP implementation does not use certain message types.

When you enable message blocking for a message type in a VoIP profile, whenever a security policy containing the VoIP profile accepts a SIP message of this type, the SIP ALG silently discards the message and records a log message about the action.

Use the following command to configure a VoIP profile to block SIP CANCEL and Update request messages:

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set block-cancel enable
      set block-update enable
    end
  end
```

SIP uses a variety of text-based messages or requests to communicate information about SIP clients and servers to the various components of the SIP network. Since SIP requests are simple text messages and since the requests or their replies can contain information about network components on either side of the FortiGate, it may be a security risk to allow these messages to pass through.

The following table lists all of the VoIP profile SIP request message blocking options. All of these options are disabled by default.



Blocking SIP OPTIONS messages may prevent a redundant configuration from operating correctly. See [Supporting geographic redundancy when blocking OPTIONS messages on page 3191](#) for information about resolving this problem.

Options for blocking SIP request messages


SIP request message	SIP message blocking CLI Option
ACK	block-ack
BYE	block-bye
Cancel	block-cancel
INFO	block-info
INVITE	block-invite

SIP request message	SIP message blocking CLI Option
Message	block-message
Notify	block-notify
Options	block-options
PRACK	block-prack
Publish	block-publish
Refer	block-refer
Register	block-register
Subscribe	block-subscribe
Update	block-update

SIP rate limiting

Configurable threshold for SIP message rates per request method. Protects SIP servers from SIP overload and DoS attacks.

SIP rate limiting



FortiGate supports rate limiting for the following types of VoIP traffic:

- Session Initiation Protocol (SIP)
- Skinny Call Control Protocol (SCCP) (most versions)
- Session Initiation Protocol for Instant Messaging and Presence Leveraging Extensions (SIMPLE).

You can use rate limiting of these VoIP protocols to protect the FortiGate and your network from SIP and SCCP Denial of Service (DoS) attacks. Rate limiting protects against SIP DoS attacks by limiting the number of SIP REGISTER and INVITE requests that the FortiGate receives per second. Rate limiting protects against SCCP DoS attacks by limiting the number of SCCP call setup messages that the FortiGate receives per minute.

You configure rate limiting for a message type by specifying a limit for the number of messages that can be received per second. The rate is limited per security policy. When VoIP rate limiting is enabled for a message type, if the a single security policy accepts more messages per second than the configured rate, the extra messages are dropped and log messages are written when the messages are dropped.

Use the following command to configure a VoIP profile to limit the number of INVITE messages accepted by each security policy that the VoIP profile is added to 100 INVITE messages a second:

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set invite-rate 100
    end
  end
```

If you are experiencing denial of service attacks from traffic using these VoIP protocols, you can enable VoIP rate limiting and limit the rates for your network. Limit the rates depending on the amount of SIP and SCCP traffic that you expect the FortiGate to be handling. You can adjust the settings if some calls are lost or if the amount of SIP or SCCP traffic is affecting FortiGate performance.

The table below lists all of the VoIP profile SIP rate limiting options. All of these options are set to 0 so are disabled by default.



Blocking SIP OPTIONS messages may prevent a redundant configuration from operating correctly. See [Supporting geographic redundancy when blocking OPTIONS messages on page 3191](#) for information about resolving this problem.

Options for SIP rate limiting

SIP request message	Rate Limiting CLI Option
ACK	ack-rate
BYE	bye-rate
Cancel	cancel-rate
INFO	info-rate
INVITE	invite-rate
Message	message-rate
Notify	notify-rate
Options	options-rate
PRACK	prack-rate
Publish	publish-rate

SIP request message	Rate Limiting CLI Option
Refer	refer-rate
Register	register-rate
Subscribe	subscribe-rate
Update	update-rate

Limiting the number of SIP dialogs accepted by a security policy

In addition to limiting the rates for receiving SIP messages, you can use the following command to limit the number of SIP dialogs (or SIP calls) that the FortiGate accepts.

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set max-dialogs 2000
    end
  end
```

This command sets the maximum number of SIP dialogs that can be open for SIP sessions accepted by any security policy that you add the VoIP profile to. The default setting of 0 does not limit the number of dialogs. You can add a limit to control the number of open dialogs and raise and lower it as required. You might want to limit the number of open dialogs for protection against SIP-based attackers opening large numbers of SIP dialogs. Every dialog takes memory and FortiGate CPU resources to process. Limiting the number of dialogs may improve the overall performance of the FortiGate. Limiting the number of dialogs will not drop calls in progress but may prevent new calls from connecting.

SIP logging

You can enable SIP logging and logging of SIP violations in a VoIP profile.

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set log-call-summary enable
      set log-violations enable
    end
  end
```

To view SIP log messages go to **Log & Report > Forward Traffic**.

Inspecting SIP over SSL/TLS (secure SIP)

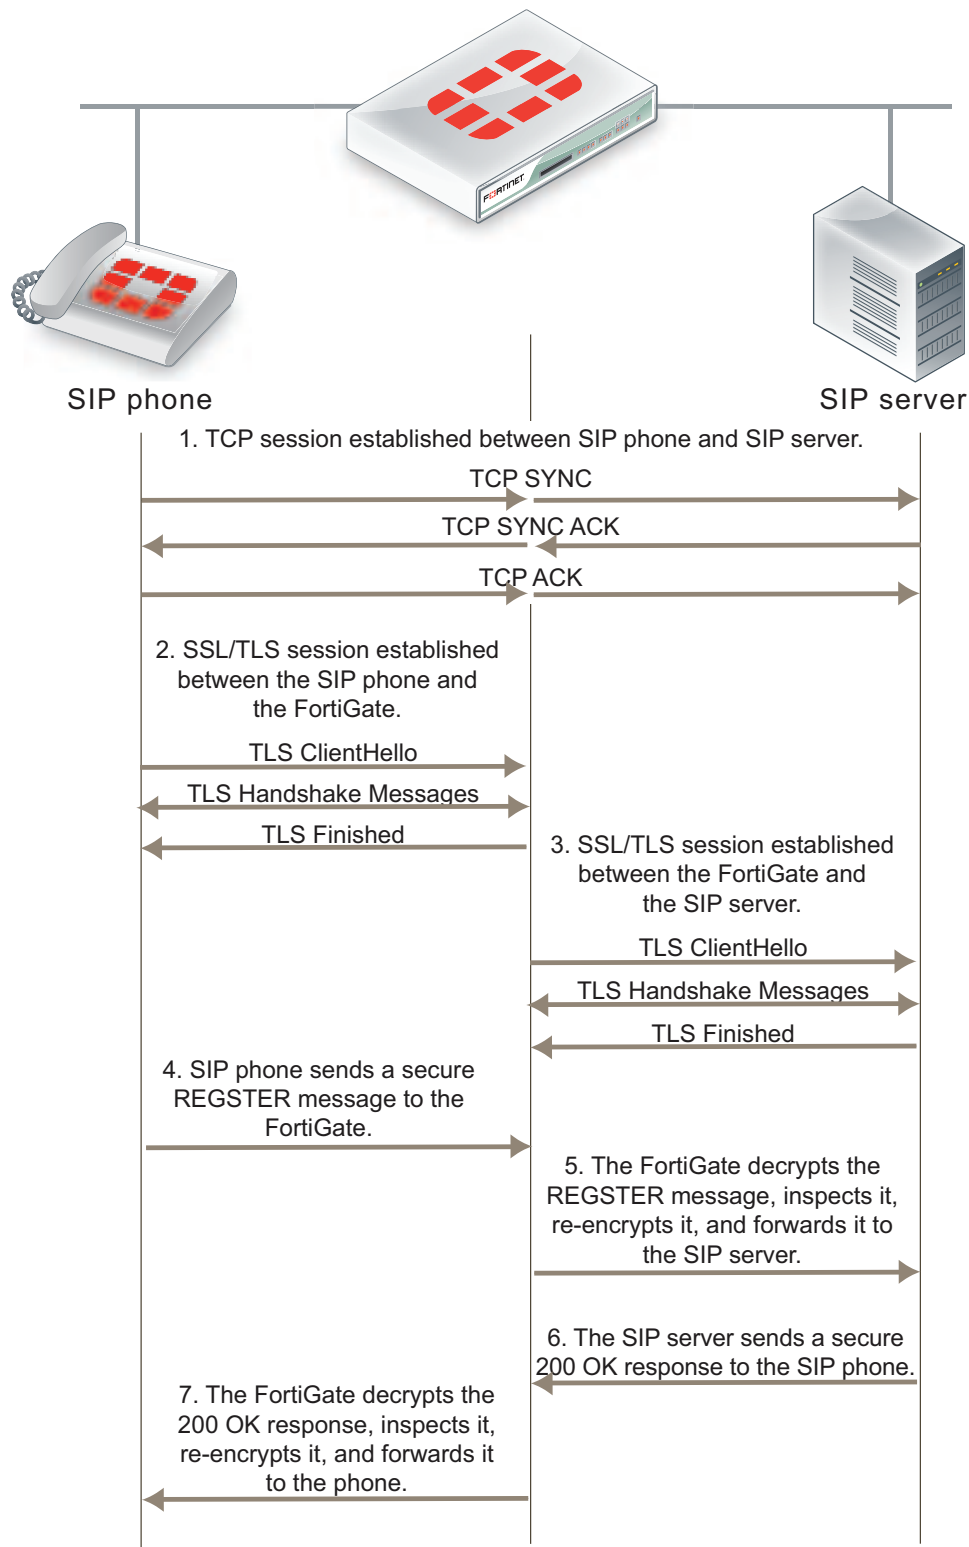
Some SIP phones and SIP servers can communicate using SSL or TLS to encrypt the SIP signaling traffic. To allow SIP over SSL/TLS calls to pass through the FortiGate, the encrypted signaling traffic has to be unencrypted and inspected. To do this, the FortiGate SIP ALG intercepts and unencrypts and inspects the SIP packets. The packets are then re-encrypted and forwarded to their destination.

Normally SIP over SSL/TLS uses port 5061. You can use the following command to change the port that the FortiGate listens on for SIP over SSL/TLS sessions to port 5066:

```
config system settings
    set sip-ssl-port 5066
end
```

The SIP ALG supports full mode SSL/TLS only. Traffic between SIP phones and the FortiGate and between the FortiGate and the SIP server is always encrypted.

You enable SSL/TLS SIP communication by enabling SSL mode in a VoIP profile. You also need to install the SIP server and client certificates on your FortiGate and add them to the SSL configuration in the VoIP profile.

SIP over SSL/TLS between a SIP phone and a SIP server

Other than enabling SSL mode and making sure the security policies accept the encrypted traffic, the FortiGate configuration for SSL/TLS SIP is the same as any SIP configuration. SIP over SSL/TLS is supported for all supported SIP configurations.

Adding the SIP server and client certificates

A VoIP profile that supports SSL/TLS SIP requires one certification for the SIP server and one certificate that is used by all of the clients. Use the following steps to add these certificates to the FortiGate. Before you start, make sure the client and server certificate files and their key files are accessible from the management computer.

1. Go to **System > Certificates** and select **Import**.
2. Set **Type** to **Certificate**.
3. Browse to the **Certificate file** and the **Key file** and select **OK**.
4. Enter a password for the certificate and select **OK**.
The certificate and key are uploaded to the FortiGate and added to the **Local Certificates** List.
5. Repeat to upload the other certificate.

The certificates are added to the list of Local Certificates as the filenames you uploaded. You can add comments to make it clear where its from and how it is intended to be used.

Adding SIP over SSL/TLS support to a VoIP profile

Use the following commands to add SIP over SSL/TLS support to the default VoIP profile. The following command enables SSL mode and adds the client and server certificates and passwords, the same ones you entered when you imported the certificates:

```
config voip profile
  edit default
    config sip
      set ssl-mode full
      set ssl-client-certificate "Client_cert"
      set ssl-server-certificate "Server_cert"
      set ssl-auth-client "check-server"
      set ssl-auth-server "check-server-group"
    end
  end
```

Other SSL mode options are also available:

<code>ssl-send-empty-frags {disable enable}</code>	Enable to send empty fragments to avoid CBC IV attacks. Compatible with SSL 3.0 and TLS 1.0 only. Default is <code>enable</code> .
<code>ssl-client-renegotiation {allow deny secure}</code>	Control how the ALG responds when a client attempts to renegotiate the SSL session. You can allow renegotiation or block sessions when the client attempts to renegotiate. You can also select <code>secure</code> to reject an SSL connection that does not support RFC 5746 secure renegotiation indication. Default is <code>allow</code> .

<code>ssl-algorithm {high low medium}</code>	Select the relative strength of the algorithms that can be selected. You can select <code>high</code> , the default, to allow only AES or 3DES, <code>medium</code> , to allow AES, 3DES, or RC4 or <code>low</code> , to allow AES, 3DES, RC4, or DES.
<code>ssl-pfs {allow deny require}</code>	Select whether to allow, deny, or require perfect forward secrecy (PFS). Default is <code>allow</code> .
<code>ssl-min-version {ssl-3.0 tls-1.0 tls-1.1}</code>	Select the minimum level of SSL support to allow. The default is <code>ssl-3.0</code> .
<code>ssl-max-version {ssl-3.0 tls-1.0 tls-1.1}</code>	Select the maximum level of SSL support to allow. The default is <code>tls-1.1</code> .

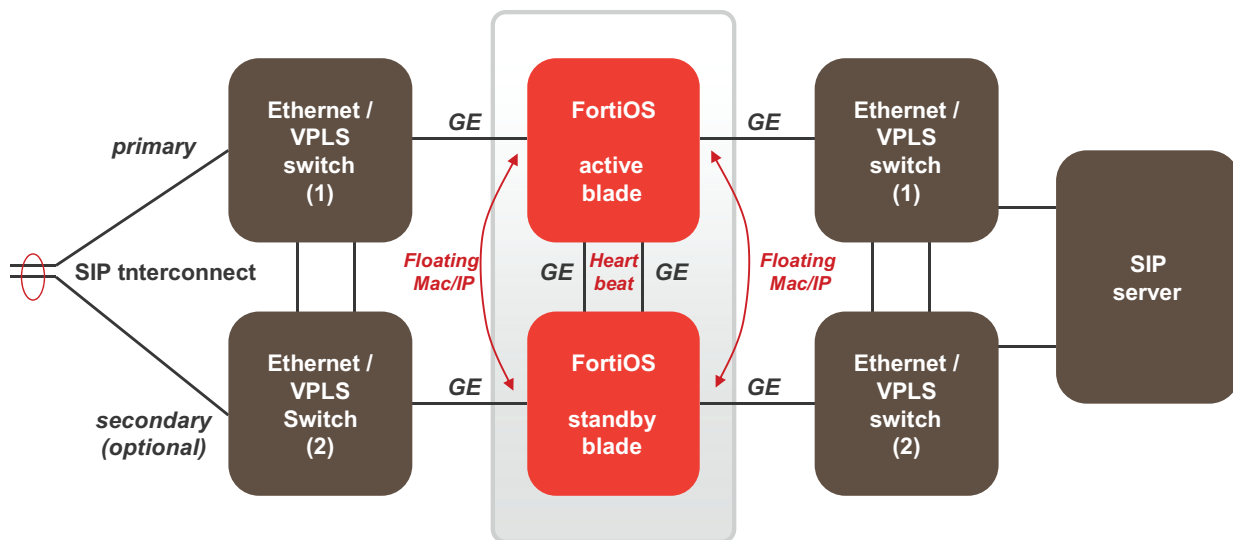
SIP and HA-session failover and geographic redundancy

FortiGate active-passive high availability (HA) supports SIP UDP session failover (also called stateful failover) if the SIP sessions are processed by the SIP ALG. To support SIP UDP session failover, create a standard HA configuration and select the enable **Session Pick-up** option.

SIP session failover replicates SIP states to all cluster units. If an HA failover occurs, all in-progress SIP UDP calls (setup complete) and their RTP flows are maintained and the calls continue after the failover with minimal or no interruption.

SIP calls being set up at the time of a failover may lose signaling messages. In most cases the SIP clients and servers should use message retransmission to complete the call setup after the failover has completed. As a result, SIP users may experience a delay if their calls are being set up when an HA failover occurs. But in most cases the call setup should be able to continue after the failover.

SIP HA session failover

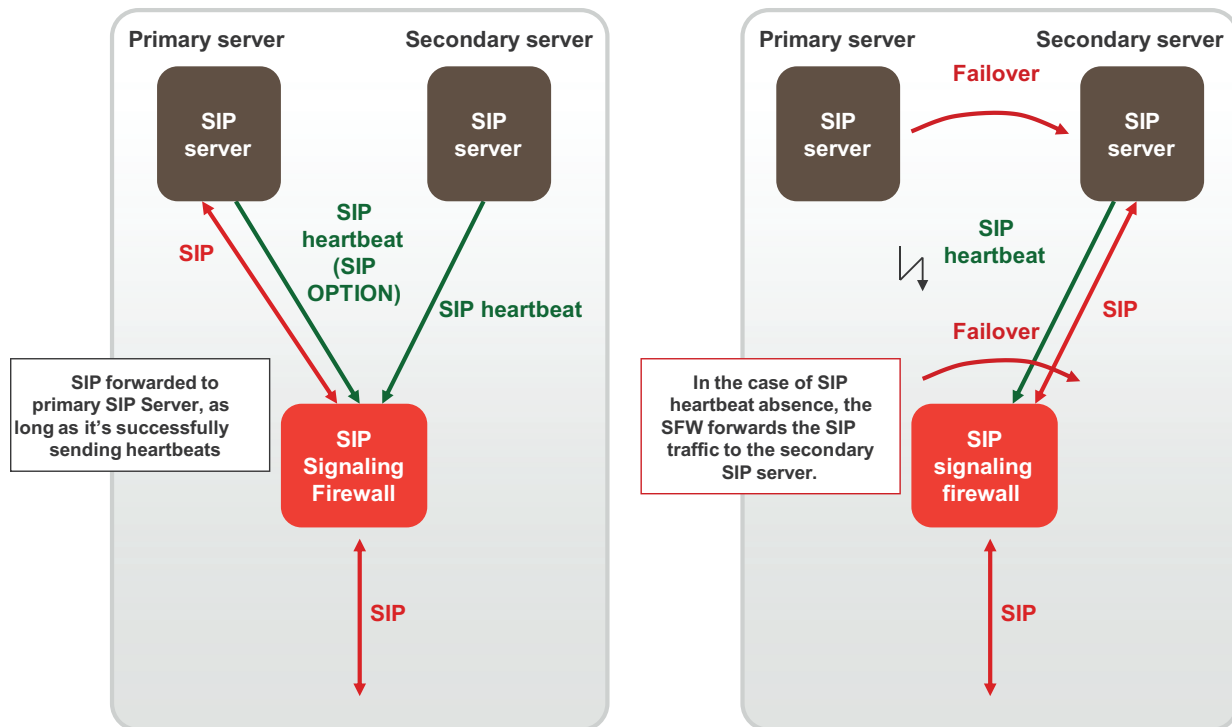


In some cases, failover during call tear down can result in hanging RTP connections that can accumulate over time and use up system memory. If this becomes a problem, you can set a time for the `call-keepalive` SIP VoIP profile setting. This setting causes the FortiGate to terminate calls with no activity after the time limit has exceeded. Range is 1 to 10,080 seconds. This options should be used with caution because it results in extra FortiGate CPU overhead and can cause delay and jitter for the VoIP call. Also, the FortiGate terminates the call without sending SIP messages to end the call. And if the SIP endpoints send SIP messages to terminate the call they will be blocked by the FortiGate if they are sent after the FortiGate terminates the call.

SIP geographic redundancy

Maintains a active-standby SIP server configuration, which even supports geographical distribution. If the active SIP server fails (missing SIP heartbeat messages or SIP traffic) FortiOS will redirect the SIP traffic to a secondary SIP server. No special configuration is required for geographic redundancy, just standard HA configuration.

Geographic redundancy



Supporting geographic redundancy when blocking OPTIONS messages

For some geographic redundant SIP configurations, the SIP servers may use SIP OPTIONS messages as heartbeats to notify the FortiGate that they are still operating (or alive). This is a kind of passive SIP monitoring mechanism where the FortiGate isn't actively monitoring the SIP servers and instead the FortiGate passively receives and analyzes OPTIONS messages from the SIP servers.

If FortiGates block SIP OPTIONS messages because `block-options` is enabled, the configuration may fail to operate correctly because the OPTIONS messages are blocked by one or more FortiGates.

However, you can work around this problem by enabling the `block-geo-red-options` application control list option. This option causes the FortiGate to refresh the local SIP server status when it receives an OPTIONS message before dropping the message. The end result is the heartbeat signals between geographically redundant SIP servers are maintained but OPTIONS messages do not pass through the FortiGate.

Use the following command to block OPTIONS messages while still supporting geographic redundancy:

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set block-options disable
      set block-geo-red-options enable
    end
  end
```



The `block-options` option setting overrides the `block-geo-red-options` option. If `block-options` is enabled the FortiGate only blocks SIP OPTIONS messages and does not refresh local SIP server status.

Support for RFC 2543-compliant branch parameters

RFC 3261 is the most recent SIP RFC, it obsoletes RFC 2543. However, some SIP implementations may use RFC 2543-compliant SIP calls.

The `rfc2543-branch` VoIP profile option allows the FortiGate to support SIP calls that include an RFC 2543-compliant branch parameter in the SIP Via header. This option also allows FortiGates to support SIP calls that include Via headers that are missing the branch parameter.

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set rfc2543-branch enable
    end
  end
```

SIP and IPS

You can enable IPS in security policies that also accept SIP sessions to protect the SIP traffic from SIP-based attacks. If you enable IPS in this way then by default the pinholes that the SIP ALG creates to allow RTP and RTCP to flow through the firewall will also have IPS enabled.

This inheritance of the IPS setting can cause performance problems if the RTP traffic volume is high since IPS checking may reduce performance in some cases. Also if you are using network processor (NP) interfaces to accelerate VoIP performance, when IPS is enabled for the pinhole traffic is diverted to the IPS and as a result is not accelerated by the network processors.

You can use the following CLI command to disable IPS for the RTP pinhole traffic.

```
config voip profile
  edit VoIP_Pro_Name
    config sip
      set ips-rtp disable
    end
  end
```


SIP debugging

This chapter includes some information to help with debugging SIP configurations.

SIP debug log format

Assuming that diagnose debug console timestamp is enabled then the following shows the debug that is generated for an INVITE if diag debug appl sip -1 is enabled:

```
2010-01-04 21:39:59 sip port 26 locate session for 192.168.2.134:5061 ->
172.16.67.192:5060
2010-01-04 21:39:59 sip sess 0x979df38 found for 192.168.2.134:5061 ->
172.16.67.192:5060
2010-01-04 21:39:59 sip port 26 192.168.2.134:5061 -> 172.16.67.192:5060
2010-01-04 21:39:59 sip port 26 read [(0,515)
(494e56495445207369703a73657276696365403139322e3136382e322e3130303a35303630205349502f322e300d0a5669613a2
05349502f322e302f554450203132372e302e312e313a353036313b6272616e63683d7a39684734624b2d363832372d3632302d3
00d0a46726f6d3a2073697070203c7369703a73697070403132372e302e312e313a353036313e3b7461673d36383237534950705
4616730303632300d0a546f3a20737574203c7369703a73657276696365403139322e3136382e322e3130303a353036303e0d0a4
3616c6c2d49443a203632302d36383237403132372e302e312e310d0a435365713a203120494e564954450d0a436f6e746163743
a207369703a73697070403132372e302e312e313a353036310d0a4d61782d466f7277617264733a2037300d0a5375626a6563743
a20506572666f726d616e636520546573740d0a436f6e74656e742d547970653a206170706c69636174696f6e2f7364700d0a436
f6e74656e742d4c656e6774683a20203132390d0a0d0a763d300d0a6f3d75736572312035333635337363520323335333638373
6333720494e20495034203132372e302e312e310d0a733d2d0d0a633d494e20495034203132372e302e312e310d0a743d3020300
d0a6d3d617564696f2036303031205254502f41565020300d0a613d7274706d61703a302050434d552f383030300d0a) (INVITE
sip:service@192.168.2.100:5060 SIP/2.0..Via: SIP/2.0/UDP
127.0.1.1:5061;branch=z9hG4bK-6827-620-0..From: sipp
%lt;sip:sipp@127.0.1.1:5061>;tag=6827SIPpTag00620..To: sut
%lt;sip:service@192.168.2.100:5060>..Call-ID: 620-6827@127.0.1.1..CSeq: 1
INVITE..Contact: sip:sipp@127.0.1.1:5061..Max-Forwards: 70..Subject: Performance
Test..Content-Type: application/sdp..Content-Length: 129....v=0..o=user1 53655765
2353687637 IN IP4 127.0.1.1..s=-..c=IN IP4 127.0.1.1..t=0 0..m=audio 6001 RTP/AVP
0..a=rtpmap:0 PCMU/8000..)]
2010-01-04 21:39:59 sip port 26 len 515
2010-01-04 21:39:59 sip port 26 INVITE '192.168.2.100:5060' addr 192.168.2.100:5060
2010-01-04 21:39:59 sip port 26 CSeq: 1 INVITE
2010-01-04 21:39:59 sip port 26 Via: UDP 127.0.1.1:5061 len 14 received 0 rport 0 0 branch 'z9hG4bK-
6827-620-0'
2010-01-04 21:39:59 sip port 26 From: 'sipp ;tag=6827SIPpTag00620' URI 'sip:sipp@127.0.1.1:5061' tag
'6827SIPpTag00620'
2010-01-04 21:39:59 sip port 26 To: 'sut ' URI 'sip:service@192.168.2.100:5060' tag ''
2010-01-04 21:39:59 sip port 26 Call-ID: '620-6827@127.0.1.1'
2010-01-04 21:39:59 sip port 26 Contact: '127.0.1.1:5061' addr 127.0.1.1:5061 expires 0
2010-01-04 21:39:59 sip port 26 Content-Length: 129 len 3
2010-01-04 21:39:59 sip port 26 sdp o=127.0.1.1 len=9
2010-01-04 21:39:59 sip port 26 sdp c=127.0.1.1 len=9
2010-01-04 21:39:59 sip port 26 sdp m=6001 len=4
2010-01-04 21:39:59 sip port 26 find call 0 '620-6827@127.0.1.1'
2010-01-04 21:39:59 sip port 26 not found
2010-01-04 21:39:59 sip port 26 call 0x97a47c0 open (collision (nil))
2010-01-04 21:39:59 sip port 26 call 0x97a47c0 open txn 0x979f7f8 INVITE dir 0
2010-01-04 21:39:59 sip port 26 sdp i: 127.0.1.1:6001
2010-01-04 21:39:59 sip port 26 policy id 1 is_client_vs_policy 1 policy_dir_rev 0
2010-01-04 21:39:59 sip port 26 policy 1 not RTP policy
2010-01-04 21:39:59 sip port 26 learn sdp from stream address
2010-01-04 21:39:59 sip port 26 call 0x97a47c0 sdp 172.16.67.198:43722
2010-01-04 21:39:59 sip port 26 call 0x97a47c0 txn 0x979f7f8 127.0.1.1:5061 find new address and port
2010-01-04 21:39:59 sip port 26 call 0x97a47c0 txn 0x979f7f8 127.0.1.1:5061 find new address and port
2010-01-04 21:39:59 sip port 26 call 0x97a47c0 txn 0x979f7f8 127.0.1.1:5061 find new address and port
2010-01-04 21:39:59 sip port 30 write 192.168.2.134:5061 -> 172.16.67.192:5060 (13,539)
2010-01-04 21:39:59 sip port 30 write [(13,539)
(494e56495445207369703a73657276696365403137322e31362e36372e3139323a35303630205349502f322e300d0a5669613a2
05349502f322e302f554450203137322e31362e36372e3139383a35323036353b6272616e63683d7a39684734624b2d363832372
```

```
d3632302d300d0a46726f6d3a2073697070203c7369703a73697070403137322e31362e36372e3139383a34333732343e3b74616
73d363832375349507054616730303632300d0a546f3a20737574203c7369703a73657276696365403137322e31362e36372e313
9323a353036303e0d0a43616c6c2d49443a203632302d36383237403132372e302e312e310d0a435365713a203120494e5649544
50d0a436f6e746163743a207369703a73697070403137322e31362e36372e3139383a34333732350d0a4d61782d466f727761726
4733a2037300d0a5375626a6563743a20506572666f726d616e636520546573740d0a436f6e74656e742d547970653a206170706
c69636174696f6e2f7364700d0a436f6e74656e742d4c656e6774683a20203133380d0a0d0a763d300d0a6f3d757365723120353
336353537363520323335336383736333720494e20495034203137322e31362e36372e3139380d0a733d2d0d0a633d494e20495
034203137322e31362e36372e3139380d0a743d3020300d0a6d3d617564696f203433373232205254502f41565020300d0a613d7
274706d61703a302050434d552f383030300d0a) (INVITE sip:service@172.16.67.192:5060 SIP/2.0..Via: SIP/2.0/UDP
172.16.67.198:52065;branch=z9hG4bK-6827-620-0..From: sipp ;tag=6827SIPpTag00620..To: sut ..Call-ID: 620-
6827@127.0.1.1..CSeq: 1 INVITE..Contact: sip:sipp@172.16.67.198:43725..Max-Forwards: 70..Subject:
Performance Test..Content-Type: application/sdp..Content-Length: 138....v=0..o=user1 53655765 2353687637
IN IP4 172.16.67.198..s=-..c=IN IP4 172.16.67.198..t=0 0..m=audio 43722 RTP/AVP 0..a=rtptime:0
PCMU/8000..)]
```

SIP-proxy filter per VDOM

You can use the `diagnose sys sip-proxy xxx` command in a VDOM to get info about how SIP is operating in each VDOM.

SIP-proxy filter command

Use the `diagnose system sip-proxy filter` to filter diagnose information for the SIP ALG. The following filters are available:

```
diag sys sip-proxy filter vd
diag sys sip-proxy filter dst-addr4
diag sys sip-proxy filter dst-addr6
diag sys sip-proxy filter dst-port
diag sys sip-proxy filter identity-policy
diag sys sip-proxy filter negate
diag sys sip-proxy filter policy
diag sys sip-proxy filter policy-type
diag sys sip-proxy filter profile-group
diag sys sip-proxy filter src-addr4
diag sys sip-proxy filter src-addr6
diag sys sip-proxy filter src-port
diag sys sip-proxy filter vd
diag sys sip-proxy filter voip-profile
```

You can clear, view and negate/invert the sense of a filter using these commands:

```
diag sys sip-proxy filter clear
diag sys sip-proxy filter list
diag sys sip-proxy filter negate
```

SIP debug setting

Control of the SIP debug output is governed by the following command

```
diagnose debug application sip <debug_level_int>
```

Where the `<debug_level_int>` is a bitmask and the individual values determine whether the listed items are logged or not. The `<debug_level_int>` can be:

- | | |
|---|--|
| 1 | Configuration changes, mainly addition/deletion/modification of virtual domains. |
| 2 | TCP connection accepts or connects, redirect creation. |

4	Create or delete a session.
16	Any IO read or write.
32	An ASCII dump of all data read or written.
64	Include HEX dump in the above output.
128	Any activity related to the use of the FortiCarrier dynamic profile feature to determine the correct profile-group to use.
256	Log summary of interesting fields in a SIP call.
1024	Any activity related to SIP geo-redundancy.
2048	Any activity related to HA syncing of SIP calls.

Display SIP rate-limit data

You can use the `diagnose sys sip-proxy meters` command to display SIP rate limiting data.

For the following command output `rate 1` shows that the current (over last second) measured rate for INVITE/ACK and BYE was 1 per second, the `peak 1` shows that the peak rate recorded is 1 per second, the `max 0` shows that there is no maximum limit set, the `count 18` indicates that 18 messages were received and `drop 0` indicates that none were dropped due to being over the limit.

```

diagnose sys sip-proxy meters
sip
sip vd: 0
sip policy: 1
sip identity-policy: 0
sip policy-type: IPv4
sip profile-group:
sip dialogs: 18
sip dialog-limit: 0
sip UNKNOWN: rate 0 peak 0 max 0 count 0 drop 0
sip ACK: rate 1 peak 1 max 0 count 18 drop 0
sip BYE: rate 1 peak 1 max 0 count 18 drop 0
sip CANCEL: rate 0 peak 0 max 0 count 0 drop 0
sip INFO: rate 0 peak 0 max 0 count 0 drop 0
sip INVITE: rate 1 peak 1 max 0 count 18 drop 0
sip MESSAGE: rate 0 peak 0 max 0 count 0 drop 0
sip NOTIFY: rate 0 peak 0 max 0 count 0 drop 0
sip OPTIONS: rate 0 peak 0 max 0 count 0 drop 0
sip PRACK: rate 0 peak 0 max 0 count 0 drop 0
sip PUBLISH: rate 0 peak 0 max 0 count 0 drop 0
sip REFER: rate 0 peak 0 max 0 count 0 drop 0
sip REGISTER: rate 0 peak 0 max 0 count 0 drop 0
sip SUBSCRIBE: rate 0 peak 0 max 0 count 0 drop 0
sip UPDATE: rate 0 peak 0 max 0 count 0 drop 0
sip PING: rate 0 peak 0 max 0 count 0 drop 0
sip YAHOOREF: rate 0 peak 0 max 0 count 0 drop 0

```

Supported RFCs

This document lists the RFCs that FortiOS supports.

What's new in FortiOS 6.0.1

FortiOS 6.0.1 introduces support for the following RFCs:

- [RFC 4389](#): Neighbor Discovery Proxies (ND Proxy)
- [RFC 5282](#): Using Authenticated Encryption Algorithms with the Encrypted Payload of the Internet Key Exchange version 2 (IKEv2) Protocol

What's new in FortiOS 6.0

FortiOS 6.0 introduces support for the following RFCs:

- [RFC 8031](#): Curve25519 and Curve448 for the Internet Key Exchange Protocol Version 2 (IKEv2) Key Agreement
- [RFC 7634](#): ChaCha20, Poly1305, and Their Use in the Internet Key Exchange Protocol (IKE) and IPsec
- [RFC 5425](#): Transport Layer Security (TLS) Transport Mapping for Syslog

Supported RFCs

FortiOS supports the following RFCs.

BGP

- [RFC 4724](#): Graceful Restart Mechanism for BGP
- [RFC 4456](#): BGP Route Reflection: An Alternative to Full Mesh Internal BGP (IBGP)
- [RFC 4360](#): BGP Extended Communities Attribute
- [RFC 4271](#): A Border Gateway Protocol 4 (BGP-4)
- [RFC 2918](#): Route Refresh Capability for BGP-4
- [RFC 2545](#): Use of BGP-4 Multiprotocol Extensions for IPv6 Inter-Domain Routing
- [RFC 2439](#): BGP Route Flap Damping
- [RFC 1997](#): BGP Communities Attribute
- [RFC 1930](#): Guidelines for creation, selection, and registration of an Autonomous System (AS)
- [RFC 1772](#): Application of the Border Gateway Protocol in the Internet

Cryptography

- [RFC 8031](#): Curve25519 and Curve448 for the Internet Key Exchange Protocol Version 2 (IKEv2) Key Agreement
- [RFC 7634](#): ChaCha20, Poly1305, and Their Use in the Internet Key Exchange Protocol (IKE) and IPsec
- [RFC 7627](#): Transport Layer Security (TLS) Session Hash and Extended Master Secret Extension
- [RFC 7539](#): ChaCha20 and Poly1305 for IETF Protocols
- [RFC 7427](#): Signature Authentication in the Internet Key Exchange Version 2 (IKEv2)
- [RFC 7383](#): Internet Key Exchange Protocol Version 2 (IKEv2) Message Fragmentation
- [RFC 7296](#): Internet Key Exchange Protocol Version 2 (IKEv2)
- [RFC 7027](#): Elliptic Curve Cryptography (ECC) Brainpool Curves for Transport Layer Security (TLS)
- [RFC 6989](#): Additional Diffie-Hellman Tests for the Internet Key Exchange Protocol Version 2 (IKEv2)
- [RFC 6954](#): Using the Elliptic Curve Cryptography (ECC) Brainpool Curves for the Internet Key Exchange Protocol Version 2 (IKEv2)
- [RFC 6290](#): A Quick Crash Detection Method for the Internet Key Exchange Protocol (IKE)
- [RFC 6023](#): A Childless Initiation of the Internet Key Exchange Version 2 (IKEv2) Security Association (SA)
- [RFC 5723](#): Internet Key Exchange Protocol Version 2 (IKEv2) Session Resumption
- [RFC 5282](#): Using Authenticated Encryption Algorithms with the Encrypted Payload of the Internet Key Exchange version 2 (IKEv2) Protocol
- [RFC 5280](#): Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile
- [RFC 4754](#): IKE and IKEv2 Authentication Using the Elliptic Curve Digital Signature Algorithm (ECDSA)
- [RFC 4635](#): HMAC SHA TSIG Algorithm Identifiers
- [RFC 4492](#): Elliptic Curve Cryptography (ECC) Cipher Suites for Transport Layer Security (TLS)
- [RFC 4478](#): Repeated Authentication in Internet Key Exchange (IKEv2) Protocol
- [RFC 4106](#): The Use of Galois/Counter Mode (GCM) in IPsec Encapsulating Security Payload (ESP)
- [RFC 3947](#): Negotiation of NAT-Traversal in the IKE

- [RFC 3602](#): The AES-CBC Cipher Algorithm and Its Use with IPsec
- [RFC 3526](#): More Modular Exponential (MODP) Diffie-Hellman groups for Internet Key Exchange (IKE)
- [RFC 2986](#): PKCS #10: Certification Request Syntax Specification Version 1.7
- [RFC 2845](#): Secret Key Transaction Authentication for DNS (TSIG)
- [RFC 2631](#): Diffie-Hellman Key Agreement Method
- [RFC 2451](#): The ESP CBC-Mode Cipher Algorithms
- [RFC 2410](#): The NULL Encryption Algorithm and Its Use With IPsec
- [RFC 2405](#): The ESP DES-CBC Cipher Algorithm With Explicit IV
- [RFC 2404](#): The Use of HMAC-SHA-1-96 within ESP and AH
- [RFC 2403](#): The Use of HMAC-MD5-96 within ESP and AH
- [RFC 2315](#): PKCS #7: Cryptographic Message Syntax Version 1.5
- [RFC 2104](#): HMAC: Keyed-Hashing for Message Authentication
- [RFC 2085](#): HMAC-MD5 IP Authentication with Replay Prevention
- [RFC 1422](#): Privacy Enhancement for Internet Electronic Mail: Part II: Certificate-Based Key Management
- [RFC 1321](#): The MD5 Message-Digest Algorithm
- [PKCS #12](#): PKCS 12 v1: Personal Information Exchange Syntax

DHCP

- [RFC 4361](#): Node-specific Client Identifiers for Dynamic Host Configuration Protocol Version Four (DHCPv4)
- [RFC 3736](#): Stateless Dynamic Host Configuration Protocol (DHCP) Service for IPv6
- [RFC 3633](#): IPv6 Prefix Options for Dynamic Host Configuration Protocol (DHCP) version 6
- [RFC 3456](#): Dynamic Host Configuration Protocol (DHCPv4) Configuration of IPsec Tunnel Mode
- [RFC 3315](#): Dynamic Host Configuration Protocol for IPv6 (DHCPv6)
- [RFC 2132](#): DHCP Options and BOOTP Vendor Extensions
- [RFC 2131](#): Dynamic Host Configuration Protocol

Diffserv

- [RFC 3260](#): New Terminology and Clarifications for Diffserv
- [RFC 2597](#): Assured Forwarding PHB Group
- [RFC 2475](#): An Architecture for Differentiated Services
- [RFC 2474](#): Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers

DNS

- [RFC 6895](#): Domain Name System (DNS) IANA Considerations
- [RFC 6604](#): xNAME RCODE and Status Bits Clarification
- [RFC 6147](#): DNS64: DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers
- [RFC 4592](#): The Role of Wildcards in the Domain Name System
- [RFC 4035](#): Protocol Modifications for the DNS Security Extensions
- [RFC 4034](#): Resource Records for the DNS Security Extensions
- [RFC 4033](#): DNS Security Introduction and Requirements
- [RFC 3597](#): Handling of Unknown DNS Resource Record (RR) Types

- [RFC 3226](#): DNSSEC and IPv6 A6 aware server/resolver message size requirements
- [RFC 3007](#): Secure Domain Name System (DNS) Dynamic Update
- [RFC 2308](#): Negative Caching of DNS Queries (DNS NCACHE)
- [RFC 2181](#): Clarifications to the DNS Specification
- [RFC 2136](#): Dynamic Updates in the Domain Name System (DNS UPDATE)
- [RFC 1996](#): A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY)
- [RFC 1995](#): Incremental Zone Transfer in DNS
- [RFC 1982](#): Serial Number Arithmetic
- [RFC 1876](#): A Means for Expressing Location Information in the Domain Name System
- [RFC 1706](#): DNS NSAP Resource Records
- [RFC 1183](#): New DNS RR Definitions
- [RFC 1101](#): DNS Encoding of Network Names and Other Types
- [RFC 1035](#): Domain Names - Implementation and Specification
- [RFC 1034](#): Domain Names - Concepts and Facilities

ICMP

- [RFC 6918](#): Formally Deprecating Some ICMPv4 Message Types
- [RFC 6633](#): Deprecation of ICMP Source Quench Messages
- [RFC 4884](#): Extended ICMP to Support Multi-Part Messages
- [RFC 4443](#): Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification
- [RFC 1191](#): Path MTU Discovery
- [RFC 792](#): Internet Control Message Protocol

IP

- [RFC 5798](#): Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6
- [RFC 4301](#): Security Architecture for the Internet Protocol
- [RFC 3272](#): Overview and Principles of Internet Traffic Engineering
- [RFC 3168](#): The Addition of Explicit Congestion Notification (ECN) to IP
- [RFC 2072](#): Router Renumbering Guide
- [RFC 2071](#): Network Renumbering Overview: Why would I want it and what is it anyway?
- [RFC 1918](#): Address Allocation for Private Internets
- [RFC 1123](#): Requirements for Internet Hosts -- Application and Support
- [RFC 1122](#): Requirements for Internet Hosts -- Communication Layers
- [RFC 791](#): Internet Protocol

IP multicast

- [RFC 5059](#): Bootstrap Router (BSR) Mechanism for Protocol Independent Multicast (PIM)
- [RFC 4604](#): Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast
- [RFC 3973](#): Protocol Independent Multicast - Dense Mode (PIM-DM): Protocol Specification (Revised)
- [RFC 3956](#): Embedding the Rendezvous Point (RP) Address in an IPv6 Multicast Address

- [RFC 3306](#): Unicast-Prefix-based IPv6 Multicast Addresses
- [RFC 2365](#): Administratively Scoped IP Multicast
- [RFC 1112](#): Host Extensions for IP Multicasting

IPsec

- [RFC 4304](#): Extended Sequence Number (ESN) Addendum to IPsec Domain of Interpretation (DOI) for Internet Security Association and Key Management Protocol (ISAKMP)
- [RFC 4303](#): IP Encapsulating Security Payload (ESP)
- [RFC 3706](#): A Traffic-Based Method of Detecting Dead Internet Key Exchange (IKE) Peers

IPv4

- [RFC 6864](#): Updated Specification of the IPv4 ID Field
- [RFC 5177](#): Network Mobility (NEMO) Extensions for Mobile IPv4
- [RFC 4632](#): Classless Inter-domain Routing (CIDR): The Internet Address Assignment and Aggregation Plan
- [RFC 3927](#): Dynamic Configuration of IPv4 Link-Local Addresses
- [RFC 3021](#): Using 31-Bit Prefixes on IPv4 Point-to-Point Links
- [RFC 1812](#): Requirements for IP Version 4 Routers

IPv6

- [RFC 6343](#): Advisory Guidelines for 6to4 Deployment
- [RFC 5175](#): IPv6 Router Advertisement Flags Option
- [RFC 5095](#): Deprecation of Type 0 Routing Headers in IPv6
- [RFC 4941](#): Privacy Extensions for Stateless Address Autoconfiguration in IPv6
- [RFC 4862](#): IPv6 Stateless Address Autoconfiguration
- [RFC 4861](#): Neighbor Discovery for IP version 6 (IPv6)
- [RFC 4389](#): Neighbor Discovery Proxies (ND Proxy)
- [RFC 4213](#): Basic Transition Mechanisms for IPv6 Hosts and Routers
- [RFC 4193](#): Unique Local IPv6 Unicast Addresses
- [RFC 4007](#): IPv6 Scoped Address Architecture
- [RFC 3971](#): SEcure Neighbor Discovery (SEND)
- [RFC 3596](#): DNS Extensions to Support IP Version 6
- [RFC 3587](#): IPv6 Global Unicast Address Format
- [RFC 3493](#): Basic Socket Interface Extensions for IPv6
- [RFC 3056](#): Connection of IPv6 Domains via IPv4 Clouds
- [RFC 3053](#): IPv6 Tunnel Broker
- [RFC 2894](#): Router Renumbering for IPv6
- [RFC 2675](#): IPv6 Jumbograms
- [RFC 2185](#): Routing Aspects Of IPv6 Transition
- [RFC 1752](#): The Recommendation for the IP Next Generation Protocol

IS-IS

- [RFC 5310](#): IS-IS Generic Cryptographic Authentication
- [RFC 5308](#): Routing IPv6 with IS-IS
- [RFC 3359](#): Reserved Type, Length and Value (TLV) Codepoints in Intermediate System to Intermediate System
- [RFC 1195](#): Use of OSI IS-IS for Routing in TCP/IP and Dual Environments

LDAP

- [RFC 4513](#): Lightweight Directory Access Protocol (LDAP): Authentication Methods and Security Mechanisms
- [RFC 4512](#): Lightweight Directory Access Protocol (LDAP): Directory Information Models
- [RFC 4511](#): Lightweight Directory Access Protocol (LDAP): The Protocol
- [RFC 3494](#): Lightweight Directory Access Protocol version 2 (LDAPv2) to Historic Status

MPLS

- [RFC 7026](#): Retiring TLVs from the Associated Channel Header of the MPLS Generic Associated Channel
- [RFC 6426](#): MPLS On-Demand Connectivity Verification and Route Tracing
- [RFC 6425](#): Detecting Data-Plane Failures in Point-to-Multipoint MPLS - Extensions to LSP Ping
- [RFC 6423](#): Using the Generic Associated Channel Label for Pseudowire in the MPLS Transport Profile (MPLS-TP)
- [RFC 5586](#): MPLS Generic Associated Channel
- [RFC 5462](#): Multiprotocol Label Switching (MPLS) Label Stack Entry: "EXP" Field Renamed to "Traffic Class" Field
- [RFC 5332](#): MPLS Multicast Encapsulations
- [RFC 5129](#): Explicit Congestion Marking in MPLS
- [RFC 4448](#): Encapsulation Methods for Transport of Ethernet over MPLS Networks
- [RFC 4182](#): Removing a Restriction on the use of MPLS Explicit NULL
- [RFC 3564](#): Requirements for Support of Differentiated Services-aware MPLS Traffic Engineering
- [RFC 3469](#): Framework for Multi-Protocol Label Switching (MPLS)-based Recovery
- [RFC 3443](#): Time To Live (TTL) Processing in Multi-Protocol Label Switching (MPLS) Networks
- [RFC 3270](#): Multi-Protocol Label Switching (MPLS) Support of Differentiated Services
- [RFC 3032](#): MPLS Label Stack Encoding

NAT

- [RFC 6888](#): Common Requirements for Carrier-Grade NATs (CGNs)
- [RFC 6146](#): Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers
- [RFC 4966](#): Reasons to Move the Network Address Translator - Protocol Translator (NAT-PT) to Historic Status
- [RFC 4787](#): Network Address Translation (NAT) Behavioral Requirements for Unicast UDP
- [RFC 4380](#): Teredo: Tunneling IPv6 over UDP through Network Address Translations (NATs)
- [RFC 3948](#): UDP Encapsulation of IPsec ESP Packets
- [RFC 3022](#): Traditional IP Network Address Translator (Traditional NAT)

OSPF

- [RFC 6860](#): Hiding Transit-Only Networks in OSPF
- [RFC 6845](#): OSPF Hybrid Broadcast and Point-to-Multipoint Interface Type
- [RFC 5340](#): OSPF for IPv6
- [RFC 4812](#): OSPF Restart Signaling
- [RFC 4811](#): OSPF Out-of-Band Link State Database (LSDB) Resynchronization
- [RFC 4203](#): OSPF Extensions in Support of Generalized Multi-Protocol Label Switching (GMPLS)
- [RFC 3630](#): Traffic Engineering (TE) Extensions to OSPF Version 2
- [RFC 3623](#): Graceful OSPF Restart
- [RFC 3509](#): Alternative Implementations of OSPF Area Border Routers
- [RFC 3101](#): The OSPF Not-So-Stubby Area (NSSA) Option
- [RFC 2328](#): OSPF Version 2
- [RFC 1765](#): OSPF Database Overflow
- [RFC 1370](#): Applicability Statement for OSPF

PPP

- [RFC 2516](#): A Method for Transmitting PPP Over Ethernet (PPPoE)
- [RFC 2364](#): PPP Over AAL5
- [RFC 1661](#): The Point-to-Point Protocol (PPP)

RADIUS

- [RFC 5176](#): Dynamic Authorization Extensions to Remote Authentication Dial In User Service (RADIUS)
- [RFC 2866](#): RADIUS Accounting
- [RFC 2548](#): Microsoft Vendor-specific RADIUS Attributes

RIP

- [RFC 4822](#): RIPv2 Cryptographic Authentication
- [RFC 2453](#): RIP Version 2
- [RFC 2080](#): RIPv2 for IPv6
- [RFC 1724](#): RIP Version 2 MIB Extension
- [RFC 1058](#): Routing Information Protocol

SIP

- [RFC 3960](#): Early Media and Ringing Tone Generation in the Session Initiation Protocol (SIP)
- [RFC 3325](#): Private Extensions to the Session Initiation Protocol (SIP) for Asserted Identity within Trusted Networks
- [RFC 3262](#): Reliability of Provisional Responses in the Session Initiation Protocol (SIP)
- [RFC 3261](#): SIP: Session Initiation Protocol

SNMP

- [RFC 4293](#): Management Information Base for the Internet Protocol (IP)
- [RFC 4273](#): Definitions of Managed Objects for BGP-4
- [RFC 4113](#): Management Information Base for the User Datagram Protocol (UDP)
- [RFC 4022](#): Management Information Base for the Transmission Control Protocol (TCP)
- [RFC 3635](#): Definitions of Managed Objects for the Ethernet-like Interface Types
- [RFC 3417](#): Transport Mappings for the Simple Network Management Protocol (SNMP)
- [RFC 3416](#): Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)
- [RFC 3414](#): User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)
- [RFC 3413](#): Simple Network Management Protocol (SNMP) Applications
- [RFC 3412](#): Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)
- [RFC 3411](#): An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks
- [RFC 3410](#): Introduction and Applicability Statements for Internet Standard Management Framework
- [RFC 2863](#): The Interfaces Group MIB
- [RFC 2578](#): Structure of Management Information Version 2 (SMIv2)
- [RFC 1238](#): CLNS MIB for use with Connectionless Network Protocol (ISO 8473) and End System to Intermediate System (ISO 9542)
- [RFC 1215](#): A Convention for Defining Traps for use with the SNMP
- [RFC 1213](#): Management Information Base for Network Management of TCP/IP-based internets: MIB-II
- [RFC 1212](#): Concise MIB Definitions
- [RFC 1157](#): A Simple Network Management Protocol (SNMP)
- [RFC 1156](#): Management Information Base for Network Management of TCP/IP-based internets
- [RFC 1155](#): Structure and Identification of Management Information for TCP/IP-based Internets

SSL

- [RFC 6176](#): Prohibiting Secure Sockets Layer (SSL) Version 2.0
- [RFC 6101](#): The Secure Sockets Layer (SSL) Protocol Version 3.0

TCP

- [RFC 6691](#): TCP Options and Maximum Segment Size (MSS)
- [RFC 6298](#): Computing TCP's Retransmission Timer
- [RFC 6093](#): On the Implementation of the TCP Urgent Mechanism
- [RFC 793](#): Transmission Control Protocol

TLS

- [RFC 6347](#): Datagram Transport Layer Security Version 1.2
- [RFC 6066](#): Transport Layer Security (TLS) Extensions: Extension Definitions
- [RFC 5746](#): Transport Layer Security (TLS) Renegotiation Indication Extension
- [RFC 5425](#): Transport Layer Security (TLS) Transport Mapping for Syslog

- [RFC 5246](#): The Transport Layer Security (TLS) Protocol Version 1.2
- [RFC 4681](#): TLS User Mapping Extension
- [RFC 4680](#): TLS Handshake Message for Supplemental Data

VPN

- [RFC 4761](#): Virtual Private LAN Service (VPLS) Using BGP for Auto-Discovery and Signaling
- [RFC 4684](#): Constrained Route Distribution for Border Gateway Protocol/MultiProtocol Label Switching (BGP/MPLS) Internet Protocol (IP) Virtual Private Networks (VPNs)
- [RFC 4577](#): OSPF as the Provider/Customer Edge Protocol for BGP/MPLS IP Virtual Private Networks (VPNs)
- [RFC 4364](#): BGP/MPLS IP Virtual Private Networks (VPNs)
- [RFC 3715](#): IPsec-Network Address Translation (NAT) Compatibility Requirements

Other protocols

- [RFC 5357](#): A Two-Way Active Measurement Protocol (TWAMP)
- [RFC 5214](#): Intra-Site Automatic Tunnel Addressing Protocol (ISATAP)
- [RFC 4960](#): Stream Control Transmission Protocol
- [RFC 4251](#): The Secure Shell (SSH) Protocol Architecture
- [RFC 3435](#): Media Gateway Control Protocol (MGCP) Version 1.0
- [RFC 3376](#): Internet Group Management Protocol, Version 3
- [RFC 2890](#): Key and Sequence Number Extensions to GRE
- [RFC 2784](#): Generic Routing Encapsulation (GRE)
- [RFC 2661](#): Layer Two Tunneling Protocol "L2TP"
- [RFC 2637](#): Point-to-Point Tunneling Protocol (PPTP)
- [RFC 2412](#): The OAKLEY Key Determination Protocol
- [RFC 2225](#): Classical IP and ARP over ATM
- [RFC 2033](#): Local Mail Transfer Protocol
- [RFC 1413](#): Identification Protocol
- [RFC 1011](#): Official Internet Protocols
- [RFC 862](#): Echo Protocol
- [RFC 768](#): User Datagram Protocol
- [The TACACS+ Protocol](#)

Miscellaneous

- [RFC 7348](#): Virtual eXtensible Local Area Network (VXLAN): A Framework for Overlaying Virtualized Layer 2 Networks over Layer 3 Networks
- [RFC 4784](#): Verizon Wireless Dynamic Mobile IP Key Update for cdma2000(R) Networks for cdma2000(R) Networks
- [RFC 4470](#): Minimally Covering NSEC Records and DNSSEC On-line Signing
- [RFC 3985](#): Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture
- [RFC 2979](#): Behavior of and Requirements for Internet Firewalls
- [RFC 2827](#): Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing

- [RFC 2780](#): IANA Allocation Guidelines For Values In the Internet Protocol and Related Headers
- [RFC 2647](#): Benchmarking Terminology for Firewall Performance
- [RFC 2644](#): Changing the Default for Directed Broadcasts in Routers
- [RFC 2231](#): MIME Parameter Value and Encoded Word Extensions: Character Sets, Languages, and Continuations
- [RFC 1945](#): Hypertext Transfer Protocol -- HTTP/1.0
- [RFC 950](#): Internet Standard Subnetting Procedure
- [RFC 894](#): A Standard for the Transmission of IP Datagrams over Ethernet Networks



FORTINET®



Copyright© 2018 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., in the U.S. and other jurisdictions, and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. In no event does Fortinet make any commitment related to future deliverables, features, or development, and circumstances may change such that any forward-looking statements herein are not accurate. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.



FORTINET®



Copyright© 2018 Fortinet, Inc. All rights reserved. Fortinet®, FortiGate®, FortiCare® and FortiGuard®, and certain other marks are registered trademarks of Fortinet, Inc., in the U.S. and other jurisdictions, and other Fortinet names herein may also be registered and/or common law trademarks of Fortinet. All other product or company names may be trademarks of their respective owners. Performance and other metrics contained herein were attained in internal lab tests under ideal conditions, and actual performance and other results may vary. Network variables, different network environments and other conditions may affect performance results. Nothing herein represents any binding commitment by Fortinet, and Fortinet disclaims all warranties, whether express or implied, except to the extent Fortinet enters a binding written contract, signed by Fortinet's General Counsel, with a purchaser that expressly warrants that the identified product will perform according to certain expressly-identified performance metrics and, in such event, only the specific performance metrics expressly identified in such binding written contract shall be binding on Fortinet. For absolute clarity, any such warranty will be limited to performance in the same ideal conditions as in Fortinet's internal lab tests. In no event does Fortinet make any commitment related to future deliverables, features, or development, and circumstances may change such that any forward-looking statements herein are not accurate. Fortinet disclaims in full any covenants, representations, and guarantees pursuant hereto, whether express or implied. Fortinet reserves the right to change, modify, transfer, or otherwise revise this publication without notice, and the most current version of the publication shall be applicable.